

Simultaneous Embedding: Edge Orderings, Relative Positions, Cutvertices*

Thomas Bläsius[†]

Annette Karrer[†]

Ignaz Rutter[‡]

Abstract

A simultaneous embedding (with fixed edges) of two graphs $G^{(1)}$ and $G^{(2)}$ with common graph $G = G^{(1)} \cap G^{(2)}$ is a pair of planar drawings of $G^{(1)}$ and $G^{(2)}$ that coincide on G . It is an open question whether there is a polynomial-time algorithm that decides whether two graphs admit a simultaneous embedding (problem SEFE).

In this paper, we present two results. First, a set of three linear-time preprocessing algorithms that remove certain substructures from a given SEFE instance, producing a set of equivalent SEFE instances without such substructures. The structures we can remove are (1) cutvertices of the *union graph* $G^{\cup} = G^{(1)} \cup G^{(2)}$, (2) most separating pairs of G^{\cup} , and (3) connected components of G that are biconnected but not a cycle.

Second, we give an $O(n^3)$ -time algorithm solving SEFE for instances with the following restriction. Let u be a pole of a P-node μ in the SPQR-tree of a block of $G^{(1)}$ or $G^{(2)}$. Then at most three virtual edges of μ may contain common edges incident to u . All algorithms extend to the sunflower case, i.e., to the case of more than three graphs pairwise intersecting in the same common graph.

1 Introduction

A simultaneous embedding of two graphs $G^{(1)}$ and $G^{(2)}$ with common graph $G = G^{(1)} \cap G^{(2)}$ is a pair of planar drawings of $G^{(1)}$ and $G^{(2)}$, that coincide on G . The problem to decide whether a simultaneous embedding exists is called SEFE (Simultaneous Embedding with Fixed Edges). This definition extends to more than two graphs. For three graphs SEFE is NP-complete [11]. In the *sunflower case* it is required that every pair of input graphs has the same intersection. See [3] for a survey on SEFE and related problems.

There are two fundamental approaches to solving SEFE in the literature. The first approach is based on the characterization of Jünger and Schulz [15] stating that finding a

*Partly done within GRADR – EUROGIGA project no. 10-EuroGIGA-OP-003.

[†]Department of Theoretical Informatics, Karlsruhe Institute of Technology (KIT) blaesius@kit.edu, annette.karrer@student.kit.edu

[‡]Department of Theoretical Informatics, Karlsruhe Institute of Technology (KIT), Karlsruhe and Department of Applied Mathematics, Charles University, Prague. rutter@kit.edu. Work was supported by a fellowship within the Postdoc-Program of the German Academic Exchange Service (DAAD).

simultaneous embedding of two graphs G^{\circledcirc} and G^{\circledast} with common graph G is equivalent to finding planar embeddings of G^{\circledcirc} and G^{\circledast} that induce the same embedding on G . The second very recent approach by Schaefer [17] is based on Hanani-Tutte-style redrawing results. One tries to characterize the existence of a SEFE via the existence of drawings of the union graph G^{\cup} where no two independent edges of the same graph cross an odd number of times. The existence of such drawings can be expressed using a linear system of boolean equations.

When following the first approach, we need two things to describe the planar embedding of the common graph G . First, for each vertex v , a cyclic order of incident edges around v . Second, for every pair of connected components H and H' of G , the face f of H containing H' . We call this relationship the *relative position* of H' with respect to H . To find a simultaneous embedding, one needs to find a pair of planar embeddings that induce the same cyclic edge orderings (*consistent edge orderings*) and the same relative positions (*consistent relative positions*) on the common graph G .

Most previous results use the first approach but none of them considers both consistent edge orderings and relative positions. Most of them assume the common graph to be connected or to contain no cycles. The strongest results of this type are the two linear-time algorithms for the case that G is biconnected by Haeupler et al. [13] and by Angelini et al. [1] and a quadratic-time algorithm for the case where G^{\circledcirc} and G^{\circledast} are biconnected and G is connected [4]. In the latter result, SEFE is modeled as an instance of the problem SIMULTANEOUS PQ-ORDERING. On the other hand, there is a linear-time algorithm for SEFE if the common graph consists of disjoint cycles [7], which requires to ensure consistent relative positions but makes edge orderings trivially consistent.

The advantage of the second approach (Hanani-Tutte) is that it implicitly handles both, consistent edge orderings and consistent relative positions, at the same time. Thus, the results by Schaefer [17] are the first that handle SEFE instances where the common graph consists of several, non-trivial connected components. He gives a polynomial-time algorithm for the cases where each connected component of the common graph is biconnected or has maximum degree 3. Although this approach is conceptionally simple, very elegant, and combines several notions of planarity within a common framework, it has two disadvantages. The running time of the algorithms are quite high and the high level of abstraction makes it difficult to generalize the results.

Contribution & Outline. In this paper, we follow the first approach and show how to enforce consistent edge orderings and consistent relative positions at the same time, by combining different recent approaches, namely the algorithm by Angelini et al. [1], the result on SIMULTANEOUS PQ-ORDERING [4] for consistent edge orderings, and the result on disjoint cycles [7] for consistent relative positions. Note that the relative positions of connected components to each other are usually expressed in terms of faces (containing the respective component). This is no longer possible if the embeddings, and thus the set of faces, of connected components are not fixed. To overcome this issue, we show that these relative positions can be expressed in terms of relative positions with respect to a cycle basis. In addition to that, we are able to handle certain cutvertices of G^{\circledcirc} and G^{\circledast} .

More precisely, we classify a vertex v to be a *union cutvertex*, a *simultaneous cutvertex*, and an *exclusive cutvertex* if v is a cutvertex of G^{\cup} , of G^{\oplus} and G^{\ominus} but not of G^{\cup} , and of G^{\oplus} but not G^{\ominus} or the other way around, respectively. Similarly, we can define UNION SEPARATING PAIRS to be separating pairs in G^{\cup} . We present several preprocessing algorithms that simplify given instances of SEFE; see Section 3. Besides a very technical preprocessing step (Section 3.4), they remove union cutvertices and most (but not all) union separating pairs; see Theorem 3, and replace connected components of G that are biconnected with cycles. They run in linear time and can be applied independently. The latter algorithm together with the linear-time algorithm for disjoint cycles [7] improves the result by Schaefer [17] for instances where every connected component of G is biconnected to linear time.

In Section 5 we show how to solve instances that have *common P-node degree 3* and simultaneous cutvertices of *common degree* at most 3 in cubic time. A vertex has common degree k if it is a common vertex with degree k in G . An instance has common P-node degree k if, for each pole v of a P-node μ (of a block) of the input graphs, at most k virtual edges of μ contain common edges incident to v . This result relies heavily on the preprocessing algorithms excluding certain structures. Together with the preprocessing steps, this includes the case where every connected component of G is biconnected, has maximum degree 3, or is outerplanar with maximum degree 3 cutvertices. As before, this approach also applies to the sunflower case.

We would like to point out that the conference version of this article contained a flaw in the handling of relative positions that are determined by P-nodes. It was erroneously claimed that certain constraints could be expressed in terms of linear Equations over \mathbb{F}_2 . The issue has been fixed by an additional preprocessing step (see Section 3.4), which allows to exclude the problematic case when treating relative positions decided by P-nodes. See Section 5.4 for additional details.

2 Preliminaries

Connectivity & SPQR-trees. A graph is *connected* if there exists a path between any pair of vertices. A *separating k -set* is a set of k vertices whose removal disconnects the graph. Separating 1-sets and 2-sets are *cutvertices* and *separating pairs*, respectively. A connected graph is *biconnected* if it has no cut vertex and *triconnected* if it has no separating pair. The maximal biconnected components of a graph are called *blocks*. The *split components* with respect to a separating k -set are the maximal subgraphs that are not disconnected by removing the separating k -set.

The SPQR-tree \mathcal{T} of a biconnected graph G represents the decomposition of G along its *split pairs*, where a split pair is either a separating pair or a pair of adjacent vertices [9]. It can be computed in linear time [12].

Let $\{s, t\}$ be a split pair and let H_1 and H_2 be two subgraphs of G such that $H_1 \cup H_2 = G$ and $H_1 \cap H_2 = \{s, t\}$. Consider the tree consisting of two nodes μ_1 and μ_2 associated with the graphs $H_1 + st$ and $H_2 + st$, respectively. These graphs are called *skeletons* of the nodes μ_i , denoted by $\text{skel}(\mu_i)$, and the special edge st is a *virtual edge*. Let $\varepsilon_1 = st$

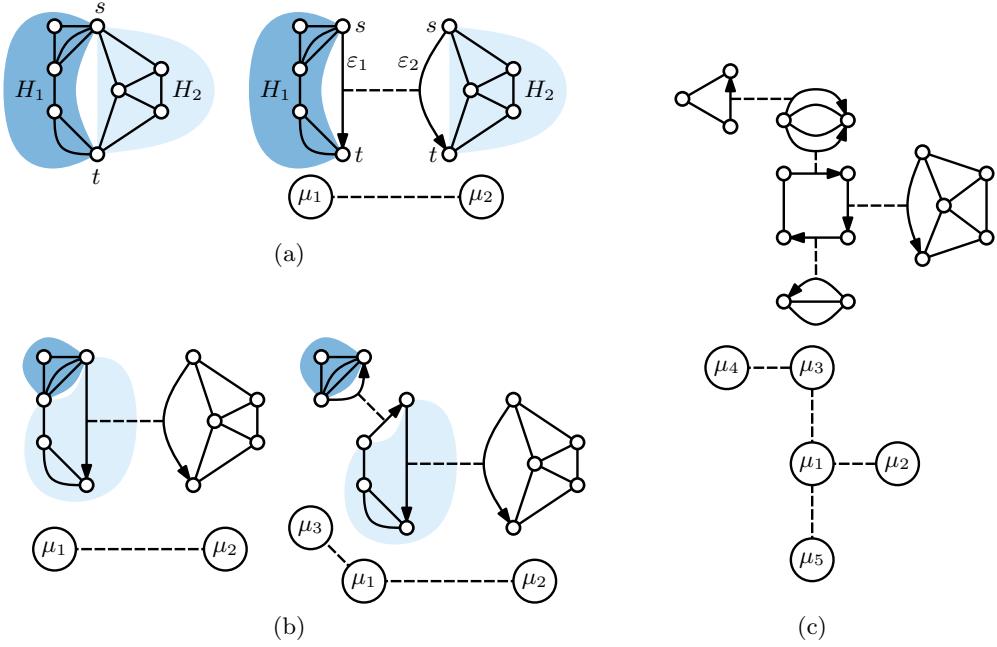


Figure 1. (a) A single decomposition step with respect to the split pair $\{s, t\}$. (b) Continued decomposition. (c) The final SPQR-tree.

and $\varepsilon_2 = st$ be the virtual edges connecting s and t in $\text{skel}(\mu_1)$ and $\text{skel}(\mu_2)$, respectively. The two edges ε_1 and ε_2 are *twins* and denote this relationship by $\text{twin}(\varepsilon_1) = \varepsilon_2$ (and vice versa). We say that the neighbor of μ_1 *corresponding* to the virtual edge ε_1 is μ_2 . Conversely, μ_1 corresponds to the virtual edge in $\text{skel}(\mu_2)$. In this way, the edge between μ_1 and μ_2 represents the twin-relationship between ε_1 and ε_2 ; see Figure 1a for an example. We can iterate this decomposition process on the graphs $\text{skel}(\mu_1)$ and $\text{skel}(\mu_2)$; see Figure 1b.

Applying this kind of decomposition systematically yields the SPQR-tree \mathcal{T} ; see Figure 1c. The skeletons of the internal nodes of \mathcal{T} are either a cycle (S-node), a bunch of parallel edges (P-node) or a triconnected planar graph (R-node). All edges in these skeletons are virtual edges. The leaves are Q-nodes and their skeleton consists of two vertices connected by a virtual and a normal edge. Sometimes it is more convenient to consider SPQR-trees without their Q-nodes. In this case, the P-, S-, and R-nodes can also contain non-virtual edges (as in Figure 1c).

When we choose a planar embedding for the skeleton of each node of the SPQR-tree \mathcal{T} , this induces a planar embedding for G . Conversely, fixing the planar embedding of G determines the embeddings of all skeletons. Thus, the combination of all planar embeddings of all skeletons is in one-to-one correspondence with the planar embeddings of G . Hence, the SPQR-tree breaks the complicated embedding choices for G (on the sphere, i.e., up to the choice of the outer face) down to embedding choices of the skeletons. These remaining choices are very simple. Skeletons of S-nodes are cycles and thus have

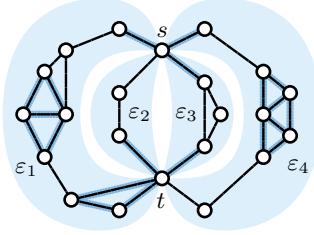


Figure 2. A P-node of G^{\circledcirc} with virtual edges $\varepsilon_1, \dots, \varepsilon_4$. The node has common P-node degree 3; for s the virtual edges $\varepsilon_1, \varepsilon_3$, and ε_4 count; for t the virtual edges $\varepsilon_1, \varepsilon_2$, and ε_3 count.

a unique planar embedding. For P-nodes we can reorder the parallel edges arbitrarily and the embedding of R-node skeletons is fixed up to a flip (i.e., up to mirroring the embedding).

Assume \mathcal{T} is rooted at an arbitrary node. In this case, the skeleton of every node (except for the root) has a unique virtual edge corresponding to its parent in \mathcal{T} . We call this virtual edge the *parent edge* and its endpoints the *poles*. We recursively define the *pertinent graph* of a node μ of \mathcal{T} . If μ is a Q-node, its pertinent graph is the non-virtual edge in $\text{skel}(\mu)$. If μ is an inner node, the pertinent graph of μ is obtained by deleting the parent edge in $\text{skel}(\mu)$ and replacing each remaining virtual edge with the pertinent graph of the corresponding child.

Let ε be a virtual edge in $\text{skel}(\mu)$ and let μ' be the corresponding neighbor of μ . The *expansion graph* $\exp(\varepsilon)$ of ε is the pertinent graph of μ' when choosing μ as root. Note that the expansion and pertinent graphs are very similar concepts. However, in most cases we use the expansion graph as it is independent of the root of the SPQR-tree (and is still defined if \mathcal{T} is unrooted). Intuitively, the expansion graph of a virtual edge is the graph that is represented by that virtual edge. Note that replacing every virtual edge in $\text{skel}(\mu)$ (for any node μ of \mathcal{T}) with its expansion graph yields the graph G . A vertex in $\exp(\varepsilon)$ is an *inner vertex* if it is not an endvertex of ε .

Let $\mathcal{T}^{\circledcirc}$ be the SPQR-tree of a block of G^{\circledcirc} in an instance of SEFE and let G be the common graph. Let further μ be a P-node of $\mathcal{T}^{\circledcirc}$. We say that μ has *common P-node degree* k if both vertices in $\text{skel}(\mu)$ are incident to common edges in the expansion graphs of at most k virtual edges (note that these can be different edges for the two vertices); see Figure 2 for an example. We say that G^{\circledcirc} has common P-node degree k if each P-node in the SPQR-tree of each block of G^{\circledcirc} has common P-node degree k . If this is the case for G^{\circledcirc} and G^{\circledast} , we say that the instance of SEFE has common P-node degree k .

We use the following conventions to make handling SPQR-trees more convenient. In most cases (as above) we explicitly name the SPQR-trees we consider (e.g., \mathcal{T} , or $\mathcal{T}^{\circledcirc}$). However, sometimes it is more convenient to write $\mathcal{T}(G)$ to denote the SPQR-tree of a given graph G . The SPQR-tree is only defined for biconnected graphs. With the SPQR-tree of a non-biconnected graph, we implicitly mean a collection of SPQR-trees, one for each block. For an S-, P-, Q-, or R-node μ of the SPQR-tree of a graph G , we also

say that μ is an S-, P-, Q-, or R-node of G , respectively. These conventions for example simplify the statement “let μ be a P-node of the SPQR-tree of a block of G ” to “let μ be a P-node of G ”.

PQ-trees. A PQ-tree, originally introduced by Booth and Lueker [8], is a tree, whose inner nodes are either P-nodes or Q-nodes (note that these P-nodes have nothing to do with the P-nodes of the SPQR-tree). The order of edges around a P-node can be chosen arbitrarily, the edges around a Q-node are fixed up to a flip. In this way, a PQ-tree represents a set of orders on its leaves. A rooted PQ-tree represents linear orders, an unrooted PQ-tree represents cyclic orders (in most cases we consider unrooted PQ-trees). Given a PQ-tree T and a subset S of its leaves, there exists another PQ-tree T' representing exactly the orders represented by T where the elements in S are consecutive. The tree T' is the *reduction* of T with respect to S . The *projection* of T to S is a PQ-tree with leaves S representing exactly the orders on S that are represented by T .

The problem SIMULTANEOUS PQ-ORDERING has several PQ-trees as input that are related by identifying some of their leaves [4]. More precisely, every instance is a directed acyclic graph, where each node is a PQ-tree, and each arc (T, T') has the property that there is an injective map from the leaves of the *child* T' to the leaves of the *parent* T . For each PQ-tree in such an instance, one wants to find an order of its leaves such that for every arc (T, T') the order chosen for the parent T is an extension of the order chosen for the child T' (with respect to the injective map). We will later use instances of SIMULTANEOUS PQ-ORDERING to express relations between orderings of edges around vertices.

3 Preprocessing Algorithms

In this section, we present several algorithms that can be used as a preprocessing of a given SEFE instance. The result is usually a set of SEFE instances that admit a solution if and only if the original instance admits one. The running time of the preprocessing algorithms is linear, and so is the total size of the equivalent set of SEFE instances. Each of the preprocessing algorithms removes certain types of structures from the instance, in particular from the common graph. Namely, we show that we can eliminate union cutvertices, simultaneous cutvertices with common-degree 3, and connected components of G that are biconnected but not a cycle. None of these algorithms introduces new cutvertices in G or increases the degree of a vertex. Thus, the preprocessing algorithms can be successively applied to a given instance, removing all the claimed structures.

Let $(G^{\circledcirc}, G^{\circledast})$ be a SEFE instance with common graph $G = G^{\circledcirc} \cap G^{\circledast}$. We can equivalently encode such an instance in terms of its *union graph* $G^{\cup} = G^{\circledcirc} \cup G^{\circledast}$, whose edges are labeled $\{1\}$, $\{2\}$, or $\{1, 2\}$, depending on whether they are contained exclusively in G^{\circledcirc} , exclusively in G^{\circledast} , or in G , respectively. Any graph with such an edge coloring can be considered as a SEFE instance. Since sometimes the coloring version is more convenient, we use these notions interchangeably throughout this section.

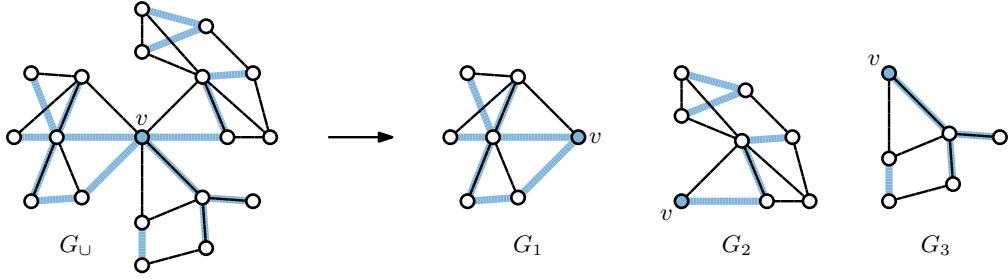


Figure 3. A union cutvertex separates a SEFE instance into independent subinstances.

3.1 Union Cutvertices

Recall that a union cutvertex of a SEFE instance $(G^{\oplus}, G^{\ominus})$ is a cutvertex of the union graph G^{\cup} . The following theorem states that the SEFE instances corresponding to the split components of a cutvertex of G^{\cup} can be solved independently; see Figure 3.

Lemma 1. *Let G^{\cup} be a SEFE instance and let v be a cutvertex of G^{\cup} with split components G_1, \dots, G_k . Then G^{\cup} admits a SEFE if and only if G_i admits a SEFE for $i = 1, \dots, k$.*

Proof. Clearly, a SEFE of G^{\cup} contains a SEFE of G_1, \dots, G_k . Conversely, given a SEFE \mathcal{E}_i of G_i for $i = 1, \dots, k$, we can assume without loss of generality that v is incident to the outer face in each of the \mathcal{E}_i . Then these embeddings can be merged to a SEFE \mathcal{E} of G^{\cup} . \square

Due to Lemma 1, it suffices to consider the blocks of G^{\cup} of a SEFE instance independently. Clearly, the blocks can be computed in $O(n)$ time, and, given a SEFE for each block, a SEFE of the original instance can be computed in $O(n)$ time.

Theorem 1. *There is a linear-time algorithm that decomposes a SEFE instance into an equivalent set of SEFE instances that do not contain union cutvertices.*

3.2 Union Separating Pairs

In analogy to a union cutvertex, we can define a *union separating pair* to be a separating pair of the union graph G^{\cup} . It is tempting to proceed as for the union cutvertices: separate G^{\cup} according to a union separating pair, solve the subinstances corresponding to the resulting subgraphs, and merge the partial solutions.

However, this approach fails as merging the partial solutions may be impossible; see Figure 4a. Note that it is easy to merge the partial solutions if all of them have u and v on the outer face of their union graph. One can enforce this kind of behaviour by connecting u and v with a common edge in each subinstance. Unfortunately, this is too restrictive as the subinstances may fail to have a SEFE with this additional edge whereas the original instance has a solution; see Figure 4b.

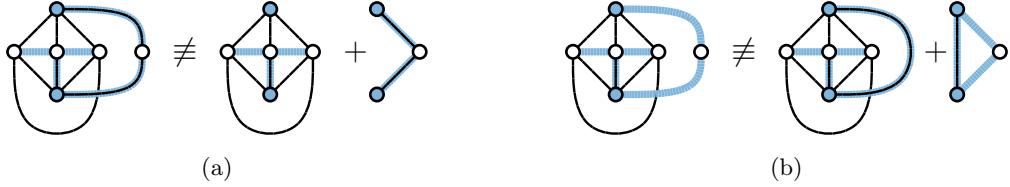


Figure 4. (a) The original instance does not admit a SEFE but the split components with respect to the separating pair $\{u, v\}$ (marked vertices) do. (b) The original instance admits a SEFE but one of the split components does not when adding the common edge uv .

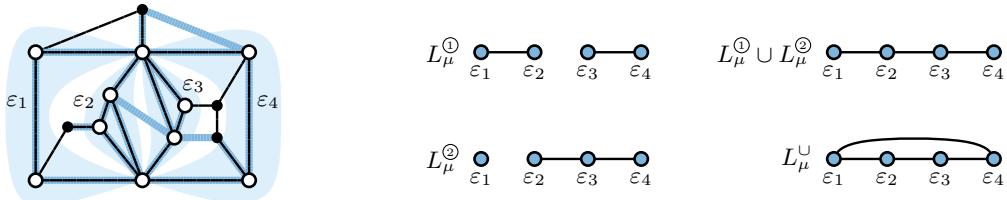


Figure 5. A P-node μ of the union graph with four virtual edges $\varepsilon_1, \dots, \varepsilon_4$ together with the link graphs L_μ^1 , L_μ^2 , $L_\mu^1 \cup L_\mu^2$, and L_μ^U .

We can, however, use the idea of adding the common edge uv in every subinstance to get rid of most union separating pairs. Throughout the whole section, we restrict our considerations to the case that u and v are vertices of the same block B of the common graph and that $\{u, v\}$ is a separating pair in B . If $\{u, v\}$ separates B into three or more split components, then u and v are poles of a P-node of $\mathcal{T}(B)$. The case when there are only two split components is a somewhat special (less interesting) case. To achieve a more concise notation, we thus assume in the following that u and v are the poles of a P-node. However, all arguments extend to the special case with two split components.

Let μ be the P-node of $\mathcal{T}(B)$ with poles u and v . Two virtual edges ε_1 and ε_2 of $\text{skel}(\mu)$ are *linked* in G^\oplus if G^\oplus contains a path from an inner vertex in $\exp(\varepsilon_1)$ to an inner vertex in $\exp(\varepsilon_2)$ that is disjoint from B (except for the end vertices of the path). The $\textcircled{1}$ -link graph L_μ^1 of μ has the virtual edges of μ as nodes, with an edge between two nodes if and only if the corresponding virtual edges are linked in G^\oplus . Analogously, we can define the $\textcircled{2}$ -link graph L_μ^2 and the union-link graph L_μ^U .

Note that the L_μ^1 and L_μ^2 are subgraphs of L_μ^U . But L_μ^U is no the union of L_μ^1 and L_μ^2 , as two virtual edges may be linked in the union graph but in none of the two exclusive graphs; see Figure 5. However, the union of L_μ^1 and L_μ^2 will also be of interest later. We call it the *exclusive-link graph* and denote it by $L_\mu^1 \cup L_\mu^2$. An edge in the exclusive-link graph indicates that the two corresponding virtual edges are either linked in G^\oplus or in G^\otimes .

We note that the following two lemmas are neither entirely new (e.g., Angelini et al. [1] use a slightly weaker statement) nor very surprising.

Lemma 2. *Let L_μ^U be a union-link graph of a given SEFE instance and let ε_1 and ε_2 be*

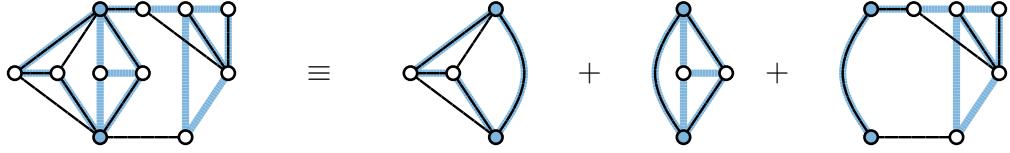


Figure 6. A union separating pair that separates a common cycle can be used to decompose the instances into simpler parts.

adjacent in L_μ^\cup . In every simultaneous embedding, ε_1 and ε_2 are adjacent in the embedding of $\text{skel}(\mu)$.

Proof. First assume that ε_1 and ε_2 are already adjacent in L_μ^\oplus . Then the expansion graphs of ε_1 and ε_2 bound a face in every embedding of G that extends to an embedding of G^\oplus . Thus, ε_1 and ε_2 must be adjacent in the embedding of $\text{skel}(\mu)$. The same holds if ε_1 and ε_2 are adjacent in G^\otimes .

Otherwise, let B be the block whose SPQR-tree contains μ . Let π be a path in the union graph connecting inner vertices v_1 and v_2 in the expansion graphs of ε_1 and ε_2 , respectively, that is disjoint from B . Clearly, then common vertices of π must be embedded into a face B that is incident to v_1 and to v_2 . Such a face only exists if ε_1 and ε_2 are adjacent in the embedding of $\text{skel}(\mu)$. \square

Lemma 3. *If (G^\oplus, G^\otimes) admits a SEFE, then each union-link graph is either a cycle or a collection of paths.*

Proof. Let B be a block of G and let μ be a P-node of $\mathcal{T}(B)$. Let $\text{skel}(\mu)$ be embedded according to a simultaneous embedding of (G^\oplus, G^\otimes) . Let $\varepsilon_1, \dots, \varepsilon_k$ be the virtual edges of $\text{skel}(\mu)$ embedded in this order. Due to Lemma 2, two virtual edges ε_i and ε_j can be adjacent in L_μ^\cup only if $i+1 = j$ or $i = k$ and $j = 1$. Thus, L_μ^\cup is a subgraph of the cycle $\varepsilon_1, \dots, \varepsilon_k, \varepsilon_1$. Hence, L_μ^\cup is either a cycle or a collection of paths. \square

Assume the union-link graph L_μ^\cup of a P-node μ is connected (i.e., by Lemma 3 a cycle or a path containing all virtual edges). Then Lemma 2 implies that the virtual edges in $\text{skel}(\mu)$ have to be embedded in a fixed order up to reversal. In this case, it remains to choose between two different embeddings, although the k virtual edges of $\text{skel}(\mu)$ have $(k-1)!$ different cyclic orders. In the following we show that we can assume without loss of generality that every union-link graph is connected.

Assume L_μ^\cup is not connected. Then the poles u and v of μ are a separating pair in the union graph. Moreover, the expansion graphs of two virtual edges from different connected components of L_μ^\cup end up in different split components with respect to u and v . Thus, we get at least two split components with a common path from u to v . If this is the case, we say that the separating pair $\{u, v\}$ separates a common cycle. We obtain the following lemma; see Figure 6.

Lemma 4. *Let $\{u, v\}$ be a separating pair of the union graph G^\cup that separates a common cycle and let $G_1^\cup, \dots, G_k^\cup$ be the split components. Then G^\cup admits a SEFE if and only if G_i^\cup with the additional common edge uv admits a SEFE for $i = 1, \dots, k$.*

Proof. Assume we have a solution for each subinstances $G_i^{\cup} + uv$. As uv is a common edge, we can assume without loss of generality that it lies in the boundary of the outer face. It is thus easy to obtain a drawing of G^{\cup} from these partial solutions without introducing any new crossings. Thus, this yields a SEFE of G^{\cup} .

Conversely, assume G^{\cup} admits a SEFE. As $\{u, v\}$ separates a common cycle, we can assume that G_1 and G_2 both contain a path of common edges connecting u and v . We have to show that $G_i + uv$ admits a SEFE for every $i = 1, \dots, k$. Assume that $i \neq 1$. Let π be the path of common edges connecting u and v in G_1 . The graph $G_i + \pi$ (which is a subgraph of G^{\cup}) admits a SEFE as the property of admitting a SEFE is closed under taking subgraphs. Moreover, it is also closed under contracting common edges. Thus, we can assume that π is actually the common edge uv . This yields a SEFE of $G_i + uv$. For $i = 1$ we can use the common path connecting u and v in G_2 instead. \square

As argued above, a disconnected union-link graph implies the existence of a separating pair that separates a common cycle. We thus obtain the following theorem.

Theorem 2. *There is a linear-time algorithm that decomposes a SEFE instance into an equivalent set of SEFE instances of total linear size in which all union-link graphs are connected.*

Proof. Clearly, applying the decomposition implied by Lemma 4 exhaustively results in a set of instances of total linear size. It remains to show that we can apply all decomposition steps in total linear time. To this end, consider the SPQR-tree \mathcal{T} of the union graph G^{\cup} . Note that G^{\cup} is non-planar in general and thus the R-nodes skeletons of \mathcal{T} may be non-planar. Nonetheless, \mathcal{T} can be computed in linear time [12] and represents all separating pairs of G^{\cup} .

Let μ be an inner node of \mathcal{T} and let $\varepsilon = uv$ be a virtual edge in $\text{skel}(\mu)$. We say that ε is a *common virtual edge* if the expansion graph of ε includes a common uv -path from. Note that $\{u, v\}$ is a separating pair of G^{\cup} . Moreover, if we know for each virtual edge whether it is a common virtual edge, we can determine whether $\{u, v\}$ separates a common cycle by only looking at $\text{skel}(\mu)$. More precisely, if μ is a P-node, then $\{u, v\}$ separates a common cycle if and only if two or more virtual edges are common virtual edges. For S- and R-nodes, $\{u, v\}$ separates a common cycle if and only if the virtual edge ε is a common virtual edge and $\text{skel}(\mu) - \varepsilon$ includes a path of common virtual edges from u to v .

Let us assume, we know for each virtual edge, whether it is a common virtual edge. Then we can easily compute the decomposition by rooting \mathcal{T} and processing it bottom up. Thus, it remains to compute the common virtual edges in linear time. To this end, first root \mathcal{T} at a Q-node. By processing \mathcal{T} bottom up, one can easily compute for each virtual edge, except for the parent edges, whether it is a common virtual edge or not.

It remains to deal with the parent edges. We process \mathcal{T} top down. When processing a node μ , we assume that we know the common virtual edges of $\text{skel}(\mu)$ (potentially including the parent edge). We then compute in $O(|\text{skel}(\mu)|)$ time for which children of μ , the parent edge is a common virtual edge. If μ is the root (i.e., a Q-node), then the only

child of μ has a common virtual edge as parent edge if and only if the edge corresponding to the Q-node μ is a common edge.

Let μ be a P-node and let ε be a virtual edge in $\text{skel}(\mu)$. Then $\text{twin}(\varepsilon)$ (which is the parent edge of the child corresponding to ε) is a common virtual edge if and only if $\text{skel}(\mu)$ includes a common virtual edge different from ε . Thus, μ can be processed in $O(|\text{skel}(\mu)|)$ time. If μ is an S-node, it similarly holds that $\text{twin}(\varepsilon)$ is a common virtual edge if and only if all virtual edges of $\text{skel}(\mu)$ except maybe ε are common virtual edges.

Finally, if μ is an R-node, consider the graph $\text{skel}'(\mu)$ obtained from $\text{skel}(\mu)$ by deleting all non-common virtual edges. Let ε be an arbitrary virtual edge of $\text{skel}(\mu)$. If ε is non-common, then $\text{twin}(\varepsilon)$ is a common virtual edge if and only if the end vertices of ε lie in the same connected component of $\text{skel}'(\mu)$. If ε is a common virtual edge, then $\text{twin}(\varepsilon)$ is a common virtual edge if and only if ε is not a bridge in $\text{skel}'(\mu)$. Note that both of these properties can be checked in constant time for each virtual edge of $\text{skel}(\mu)$ after $O(|\text{skel}(\mu)|)$ preprocessing time. Thus, we can also process R-nodes in $O(|\text{skel}(\mu)|)$ time, which yields an overall linear running time. \square

Let B be a block of the common graph and let μ be a P-node of $\mathcal{T}(B)$. By Theorem 2, we can assume that the union-link graph L_μ^\cup is connected. Thus, the ordering of the virtual edges in $\text{skel}(\mu)$ is fixed up to reversal. Hence, the embedding choices for μ are the same as those for an R-node.

In the following, we provide further simplifications by eliminating some types of simultaneous separating pairs. Let u and v be the poles of the P-node μ . Consider the case that $\{u, v\}$ is a separating pair in the union graph G^\cup with split components $G_1^\cup, \dots, G_k^\cup$ (we can assume by Theorem 1 that neither u nor v is a cutvertex in G^\cup). As before, we denote the common graph and the exclusive graphs corresponding to the SEFE instances G_i^\cup (for $i = 1, \dots, k$) by G_i , G_i^\oplus and G_i^\otimes , respectively.

We define G_i^\cup to be *common connected* if u and v are connected by a path in G_i ; see Figure 7a. The split component G_i^\cup is *exclusive connected*, if it is not common connected but u and v are connected by exclusive paths in both graphs G_i^\oplus and G_i^\otimes ; see Figure 7b. It is *\oplus -connected*, if u and v are connected by a path in G_i^\oplus but not in G_i^\otimes ; see Figure 7c. The term *\otimes -connected* is defined analogously; see Figure 7d. Note that being \oplus - or \otimes -connected excludes being common or exclusive connected. Finally, if G_i^\cup is neither of the above, it is *union connected*; see Figure 7e.

We say that μ is an *impossible P-node* if L_μ^\oplus is a cycle and one of the split components is \oplus -connected, if L_μ^\otimes is a cycle and one of the split components is \otimes -connected, or if $L_\mu^\oplus \cup L_\mu^\otimes$ is a cycle and one of the split components is exclusive connected.

Lemma 5. *A SEFE instance with an impossible P-node is a no-instance.*

Proof. Let $G_1^\cup, \dots, G_k^\cup$ be the split components with respect to the poles u and v of the P-node μ . As μ is an impossible P-node, the union-link graph L_μ^\cup is a cycle. Thus, at most one split component can be common connected. As u and v are the poles of a P-node of the common graph, one of the split components must be common connected. Thus, exactly one split component, without loss of generality G_1^\cup , is common connected.

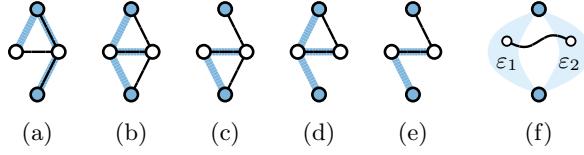


Figure 7. (a–e) A split component that is common connected, exclusive connected, ①-connected, ②-connected, and union connected, respectively. (f) The face between two virtual edges that are ①-linked (i.e., connected in $L_\mu^{\text{①}}$).

First assume that μ is an impossible P-node due to the fact that $L_\mu^{\text{①}}$ is a cycle and one of the split components, without loss of generality G_2^{U} is ①-connected.

Assume the given SEFE instance G^{U} admits a SEFE and assume G_1^{U} and G_2^{U} are embedded according to this SEFE. As G_2^{U} is ①-connected (Figure 7c), the graph $G_2^{\text{①}}$ includes a path $\pi^{\text{①}}$ from u to v . Clearly, $\pi^{\text{①}}$ lies in a single face of $G_1^{\text{①}}$. Let f be the corresponding face of the common graph G_1 . The boundary of f belongs to the expansion graphs of two different virtual edges ε_1 and ε_2 of $\text{skel}(\mu)$; see Figure 7f. However, ε_1 and ε_2 cannot be ①-linked (as in Figure 7f), as otherwise $\pi^{\text{①}}$ could not be embedded into the face f without having a crossing in $G^{\text{①}}$. It follows that $L_\mu^{\text{①}}$ cannot be a cycle, a contradiction.

Analogously, if $L_\mu^{\text{②}}$ is a cycle and one of the split components is ②-connected, we find a path $\pi^{\text{②}}$ that is a witness for a pair of adjacent virtual edges that are not ②-linked. It remains to consider the case where $L_\mu^{\text{①}} \cup L_\mu^{\text{②}}$ is a cycle and G_2^{U} is exclusive connected. In this case, $G_2^{\text{①}}$ and $G_2^{\text{②}}$ include paths $\pi^{\text{①}}$ and $\pi^{\text{②}}$, respectively, connecting u and v . As they both belong to the same split component, they have to be embedded in the same common face of G_1 . Thus, there are adjacent virtual edges that are neither ①- nor ②-linked. Hence, $L_\mu^{\text{①}} \cup L_\mu^{\text{②}}$ is not a circle. \square

Due to this lemma, it is sufficient to consider the case that μ is not an impossible P-node. We want to show that the different split components (in the union graph, with respect to the poles u and v of μ) can be handled independently. However, we have to exclude a special case to make this true. Let G_i^{U} be one of the split components that is exclusive connected. We say that G_i^{U} has *common ends* if it contains a common edge incident to u or to v . Figure 8 shows an example, where the following lemma does not hold without excluding exclusive connected components with common ends.

Lemma 6. *Let G^{U} be a SEFE instance and let μ be a non-impossible P-node whose poles are a separating pair with split components $G_1^{\text{U}}, \dots, G_k^{\text{U}}$. Assume G_1^{U} is the only common connected split component and none of the exclusive connected components has common ends. Then G^{U} admits a SEFE if and only if G_1^{U} admits a SEFE and G_i^{U} together with the common edge uv admits a SEFE for $i = 2, \dots, k$.*

Proof. Assume that G_1^{U} and $G_i^{\text{U}} + uv$ (for $i = 2, \dots, k$) admit simultaneous embeddings. We show how to combine the simultaneous embeddings of G_1^{U} and $G_2^{\text{U}} + uv$ to a simultaneous embedding of $G_1^{\text{U}} \cup G_2^{\text{U}}$. The procedure can then be iteratively applied to the other split

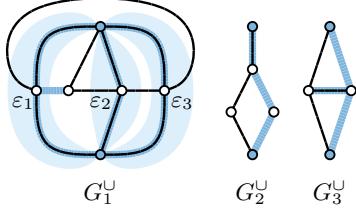


Figure 8. The union split component G_1^U includes the expansion graphs of all three virtual edges ε_1 , ε_2 , and ε_3 . The edge pairs $\varepsilon_1, \varepsilon_3$ and $\varepsilon_2, \varepsilon_3$ are ①-linked, thus the exclusive connected split component G_2^U cannot be embedded into the faces between ε_1 and ε_3 or between ε_2 and ε_3 . Although ε_1 and ε_2 are neither ①- nor ②-linked, G_2^U cannot be embedded into the face between ε_1 and ε_2 due to its common end. The component G_3^U has no common end and can be embedded into the face between ε_1 and ε_2 .

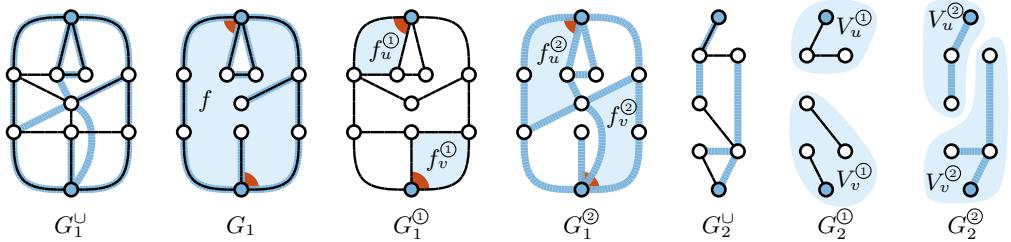


Figure 9. Two split components of the union graph illustrating the proof of Lemma 6.

components. We have to distinguish the cases that G_2^U is union connected, ①-connected, ②-connected, and exclusive connected (without common ends).

First assume that G_2^U is union connected. Figure 9 shows an example illustrating the proof for this case. As u and v are the poles of a P-node, the common graph G_1 has a face f that is incident to u and to v . Let further $f_u^①$ and $f_v^①$ be faces of $G_1^①$ incident to u and v , respectively, that are both part of the union face f . Similarly, we choose faces $f_u^②$ and $f_v^②$ in $G_1^②$ that are incident to u and v , respectively, and that are both part of f . Note that u might have several incidences to the face f , i.e., when u is a cutvertex in G_1 and one of the corresponding blocks is embedded into f . In this case, we choose $f_u^②$ such that it has *the same incidence* to u as $f_u^①$, i.e., the common edges appearing in the cyclic order around u before and after $f_u^②$ are the same as those that appear before and after $f_u^①$. We ensure the same for $f_v^①$ and $f_v^②$. In the example in Figure 9, f has two incidences to u and two incidences to v and the chosen incidence is marked by an angle.

Due to the common edge uv , we can assume that the SEFE of G_2^U has uv and uv on the outer face. As G_2^U is not ①-connected, we can separate the vertices of G_2^U into two subsets $V_u^①$ and $V_v^①$, such that $V_u^①$ contains all vertices of the connected component of G_2^U containing u , while $V_v^①$ contains all other vertices. We can then embed the vertices of $V_u^①$ into $f_u^①$ and the vertices of $V_v^①$ into $f_v^①$ without changing the embedding of G_2^U . In the same way, $G_2^②$ can be embedded into $f_u^②$ and $f_v^②$.

As we did not change the embedding of G_1^U or G_2^U , the edge orderings are consistent

for all vertices except maybe u and v . Moreover, the relative positions between connected components in G_1^{\cup} is consistent and the same holds for G_2^{\cup} . As the four faces f_u^{\circledcirc} , f_v^{\circledcirc} , f_u^{\circledast} , and f_v^{\circledast} belong to the same common face f , the relative positions of components in G_2 with respect to components in G_1 are also consistent. Moreover, all components of G_1 lie in the outer face of G_2 with respect to G_2^{\circledcirc} and G_2^{\circledast} . Finally, the edge ordering at u is consistent, as all edges incident to u in G_2^{\circledcirc} and G_2^{\circledast} are embedded between the same pair of common edges in G_1 . As the same holds for v , we obtain a simultaneous embedding of $G_1^{\cup} \cup G_2^{\cup}$.

If G_2^{\cup} is \circledcirc -connected, we know that L_{μ}^{\circledcirc} is not a circle (otherwise, μ would be impossible). Thus, we can choose the faces f , f_u^{\circledcirc} , and f_v^{\circledcirc} such that $f_u^{\circledcirc} = f_v^{\circledcirc}$. Then we can embed G_2^{\circledcirc} into this face without separating it. All remaining arguments work the same as above. The case that G_2^{\cup} is \circledast -connected is symmetric.

Finally, if G_2^{\cup} is exclusive connected, there is a pair of virtual edges that are neither \circledcirc - nor \circledast -linked. Thus, we can choose the common face f and the faces f_u^{\circledcirc} , f_v^{\circledcirc} , f_u^{\circledast} , and f_v^{\circledast} belonging to f such that $f_u^{\circledcirc} = f_v^{\circledcirc} = f^{\circledcirc}$ and $f_u^{\circledast} = f_v^{\circledast} = f^{\circledast}$. Unfortunately, we cannot always ensure that f^{\circledcirc} and f^{\circledast} have the same incidence to u or v ; see Figure 8. However, the arguments from the previous cases still ensure that all relative positions and all cyclic orders except for maybe at u and v are consistent. As G_2^{\cup} has no common ends, all common edges incident to u and v are contained in G_1 and thus the cyclic orders around these vertices are also consistent.

Note that combining the simultaneous embeddings G_1^{\cup} and G_2^{\cup} in this way (for all four cases) maintains the properties that there are faces in G_1 , G_1^{\circledcirc} , or G_2^{\circledast} that are incident to both poles u and v . Thus, we can continue adding embeddings of all remaining subinstances $G_3^{\cup}, \dots, G_k^{\cup}$ in the same way. \square

Assume we exhaustively applied Lemma 4, Lemma 5, and Lemma 6 to a given instance of SEFE and let $(G^{\circledcirc}, G^{\circledast})$ be the resulting instance. Let u and v be the poles of a P-node μ of the common graph such that $\{u, v\}$ are a separating pair in the union graph. By Lemma 4 we can assume that $\{u, v\}$ does not separate a common cycle. Thus, exactly one split component has a common uv -path. By Lemma 5, we can assume that μ is a non-impossible P-node. Thus, we could apply Lemma 6 if there were split components without common ends. Hence we obtain the following theorem.

Theorem 3. *Let $(G^{\circledcirc}, G^{\circledast})$ be an instance of SEFE. In linear time, we can find equivalent instances such that every union separating pair $\{u, v\}$ has one of the following properties.*

- *The vertices u and v are not the poles of a P-node of a common block.*
- *Every split component has a common edge incident to u or to v but only one has a common uv -path.*

Proof. It remains to prove the claimed running time. The linear running time for decomposing the instances along its union separating pairs that separate a common cycle was already shown for Theorem 2. In Section 4.3 we extend the algorithm by Angelini et al. [1] for solving SEFE if the common graph is biconnected to the case where we allow exclusive vertices and have so-called union bridge constraints. It is not hard to see that

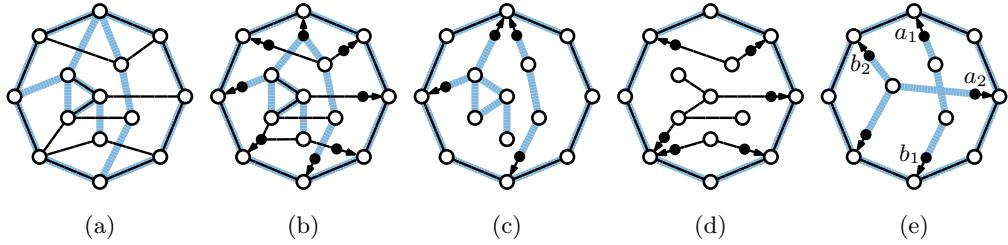


Figure 10. Situation where the connected component C of G is a cycle. (a) A simultaneous embedding of (G^1, G^2) with C on the outer face. (b) Removing C yields a single connected component in G^U . Thus, there is only one union bridge. Its attachment vertices are illustrated as black dots. (c) The two ①-bridges. (d) The three ②-bridges. Note that different bridges might share attachment vertices. (e) Two alternating ①-bridges.

testing (G^1, G^2) for the existence of impossible P-nodes can be done using the linear-time algorithm from Section 4.3.

It remains to decompose the union graph G^U according to separating pairs that separated G^U according to Lemma 6. As in the proof of Theorem 2, we consider the SPQR-tree \mathcal{T} of G^U . For Theorem 2, we had to compute for every virtual edge, whether its expansion graph included a common path between its endpoints. Now, we in addition have to know which expansion graphs are exclusive connected and have common ends. This can be done analogously to the proof of Theorem 2. \square

3.3 Connected Components that are Biconnected

Let (G^1, G^2) be a SEFE instance and let C be a connected component of the common graph G that is a cycle; see Figure 10a. A *union bridge* of G^1 and G^2 with respect to C is a connected component of $G^U - C$ together with all its *attachment vertices* on C ; see Figure 10b. Equivalently, the union bridges are the split components of G^U with respect to the vertices of C excluding the edges of C . Similarly, there are *①-bridges* and *②-bridges*, which are connected components of $G^1 - C$ and $G^2 - C$ together with their attachment vertices on C , respectively; see Figure 10c–d. We say that two bridges B_1 and B_2 *alternate* if there are attachments a_1, b_1 of B_1 and attachments a_2, b_2 of B_2 , such that the order along C is $a_1 a_2 b_1 b_2$; see Figure 10e. We have the following lemma, which basically states that we can handle different union bridges independently

Lemma 7. *Let G^1 and G^2 be two planar graphs and let C be a connected component of the common graph that is a cycle. Then the graphs G^1 and G^2 admit a SEFE where C is the boundary of the outer face if and only if (i) each union bridge admits a SEFE together with C and (ii) no two ①-bridges of C alternate for $i = 1, 2$.*

Proof. Clearly the conditions are necessary; we prove sufficiency. Let B_1, \dots, B_k be the union bridges with respect to C , and let $(\mathcal{E}_1^1, \mathcal{E}_1^2), \dots, (\mathcal{E}_k^1, \mathcal{E}_k^2)$ be the corresponding simultaneous embeddings of B_i together with C , which exist by condition (i). Note that each union bridge is connected, and hence all its edges and vertices are embedded on the

same side of C . After possibly flipping some of the embeddings, we may assume that each of them has C with the same clockwise orientation as the outer face.

We now glue $\mathcal{E}_1^{\circledcirc}, \dots, \mathcal{E}_k^{\circledcirc}$ to an embedding $\mathcal{E}^{\circledcirc}$ of G^{\circledcirc} , which is possible by condition (ii). In the same way, we find an embedding $\mathcal{E}^{\circledast}$ of G^{\circledast} from $\mathcal{E}_1^{\circledast}, \dots, \mathcal{E}_k^{\circledast}$. We claim that $(\mathcal{E}^{\circledcirc}, \mathcal{E}^{\circledast})$ is a SEFE of G^{\circledcirc} and G^{\circledast} . For the consistent edge orderings, observe that any common vertex v with common-degree at least 3 is contained, together with all neighbors, in some union bridge B_i . The compatibility of the edge ordering follows since $(\mathcal{E}_i^{\circledcirc}, \mathcal{E}_i^{\circledast})$ is a SEFE. Concerning the relative position of a vertex v and some common cycle C' , we note that the relative positions clearly coincide in $\mathcal{E}^{\circledcirc}$ and $\mathcal{E}^{\circledast}$ for $C' = C$. Otherwise C' is contained in some union bridge. If v is embedded in the interior of C' in one of the two embeddings, then it is contained in the same union bridge as C' , and the compatibility follows. If this case does not apply, it is embedded outside of C' in both embeddings, which is compatible as well. \square

We note that this approach fails, when the cycle C is not a connected component of G , i.e., when a union bridge contains common edges incident to an attachment vertex. The reason is that the order of common edges incident to this attachment vertex is chosen in the moment one reinserts the union bridges into C .

Now consider a connected component C of the common graph G of a SEFE instance such that C is biconnected. Such a component is called *2-component*. If C is a cycle, it is a *trivial 2-component*. We define the union bridges, and the ①- and ②-bridges of G^{\circledcirc} and G^{\circledast} with respect to C as above. We call an embedding \mathcal{E} of C together with an assignment of the union bridges to its faces *admissible* if and only if, (i) for each union bridge, all attachments are incident to the face to which it is assigned, and (ii) no two ①- and not two ②-bridges that are assigned to the same face alternate.

In the following, we try to solve the given SEFE instance $(G^{\circledcirc}, G^{\circledast})$ by first finding an admissible embedding of the 2-component C . Then we test for every face of C whether all union bridges can be embedded inside the corresponding facial cycle. By Lemma 7 we know that this is possible if and only if each union bridge together with the facial cycle admits a SEFE and no two ①-bridges (for $i = 1, 2$) alternate. The latter is ensured by property (ii) of the admissible embedding of C . The former yields simpler SEFE instances in which the 2-component C is represented by a simple cycle. It remains to show that, if this approach fails, there exists no SEFE. First note that the properties (i) and (ii) of an admissible embedding are clearly necessary. Thus, if there is no admissible embedding of C , then there is no SEFE. It remains to show that it does not depend on the admissible embedding of C one chooses, whether a union bridge together with the facial cycle of the face it is assigned to admits a SEFE or not. In fact, the following lemma shows that the facial cycle one gets for a union bridge is more or less independent from the embedding of C , i.e., the attachment vertices of the bridge always appear in the same order along this cycle.

Lemma 8. *Let G be a biconnected planar graph and let X be a set of vertices that are incident to a common face in some planar embedding of G . Then the order of X in any simple cycle of G containing X is unique up to reversal.*

Proof. Consider a planar embedding \mathcal{E} of G where all vertices in X share a face, and let C_X denote the corresponding facial cycle. Note that C_X is simple since G is biconnected. Let C be an arbitrary simple cycle in G containing all vertices in X . In \mathcal{E} , all parts of C that are disjoint from C_X are embedded outside of C_X . Let C'_X denote the cycle obtained from C_X by contracting all maximal paths whose internal vertices do not belong to C to single edges. Observe that C_X and C'_X visit the vertices of X in the same order. Consider the graph $C \cup C'_X$, which is clearly outerplanar and biconnected. Hence both C and C'_X visit the vertices of X in the same order (up to complete reversal). Since C was chosen arbitrarily, the claim follows. \square

For a union bridge B , let C_B denote the cycle consisting of the attachments of B in the ordering of an arbitrary cycle of G containing all the attachments. By Lemma 8, the cycle C_B is uniquely defined. Let further G_B denote the graph consisting of the union bridge B and the cycle C_B connecting the attachment vertices of B . We call this graph the *union bridge graph* of the bridge B . The following lemma formally states our above-mentioned strategy to decompose a SEFE instance.

Lemma 9. *Let G^{\circledcirc} and G^{\circledast} be two connected planar graphs and let C be a 2-component of the common graph G . Then the graphs G^{\circledcirc} and G^{\circledast} admit a SEFE if and only if (i) C admits an admissible embedding, and (ii) each union bridge graph admits a SEFE. If a SEFE exists, the embedding of C can be chosen as an arbitrary admissible embedding.*

Proof. Clearly, a SEFE of G^{\circledcirc} and G^{\circledast} defines an embedding of C and a bridge assignment that is admissible. Moreover, it induces a SEFE of each union bridge graph.

Conversely, assume that C admits an admissible embedding and each union bridge graph admits a SEFE. We obtain a SEFE of G^{\circledcirc} and G^{\circledast} as follows. Embed C with the admissible embedding and consider a face f of this embedding with facial cycle C_f . Let B_1, \dots, B_k denote the union bridges that are assigned to this face, and let $(\mathcal{E}_1^{\circledcirc}, \mathcal{E}_1^{\circledast}), \dots, (\mathcal{E}_k^{\circledcirc}, \mathcal{E}_k^{\circledast})$ be simultaneous embeddings of the bridge graphs G_B . By subdividing the cycle C_B , in each of the embeddings, we may assume that the outer face of each B_i in the embedding $(\mathcal{E}_i^{\circledcirc}, \mathcal{E}_i^{\circledast})$ is the facial cycle C_f with the same orientation in each of them. By Lemma 7, we can hence combine them to a single SEFE of all union bridges whose outer face is the cycle C_f . We embed this SEFE into the face f of C . Since we can treat the different faces of C independently, applying this step for each face yields a SEFE of G^{\circledcirc} and G^{\circledast} with the claimed embedding of C . \square

Lemma 9 suggests a simple strategy for reducing SEFE instances containing non-trivial 2-components. Namely, take such a component, construct the corresponding union bridge graphs, where C occurs only as a cycle, and find an admissible embedding of C . Finding an admissible embedding for C can be done as follows. To enforce the non-overlapping attachment property, replace each \circledcirc -bridge of C by a *dummy \circledcirc -bridge* that consists of a single vertex that is connected to the attachments of that bridge via edges in E^{\circledcirc} . Similarly, we replace \circledast -bridges, which are connected to attachments via exclusive edges in E^{\circledast} . We seek a SEFE of the resulting instance (where the common graph is biconnected), additionally requiring that dummy bridges belonging to the same union

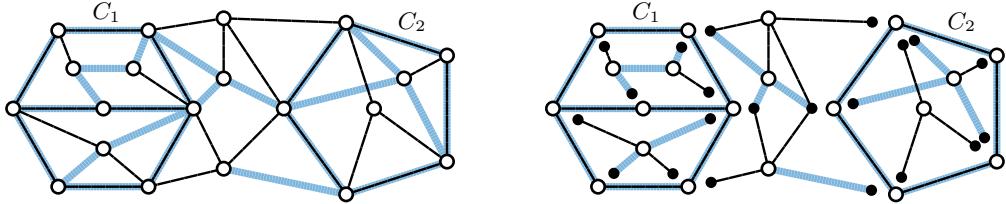


Figure 11. The instance on the left contains two 2-components C_1 and C_2 . The corresponding union subbridges are shown on the right.

bridge are embedded in the same face. We also refer to such an instance as *SEFE with union bridge constraints*. A slight modification of the algorithm by Angelini et al. [1] can decide the existence of such an embedding in polynomial time. This gives the following lemma.

Lemma 10. *Computing an admissible embedding of a 2-component C is equivalent to solving SEFE with union bridge constraints on an instance having C as common graph. This can be done in polynomial time.*

It then remains to treat the union bridge graphs. Exhaustively applying Lemma 9 (using Lemma 10 to find admissible embeddings) results in a set of SEFE instances where each 2-component is trivial. Note that we could go even further and decompose along cycles that have more than one union bridge. However, this is not necessary to obtain the following theorem.

Theorem 4. *Given a SEFE instance, an equivalent set of instances of total linear size such that each 2-component of these instances is trivial can be computed in polynomial time.*

We can improve the running time in Theorem 4 to linear. However, it is quite tedious work, involving a lot of data structures, and results in a lengthy proof. To not disturb the reading flow too much, the proof is deferred to its own section (Section 4, starting on page 20). Here, we only sketch it very roughly.

We do not apply an iterative process, removing one 2-component after another (as suggested above), but we decompose the whole instance at once. For this, we introduce the notion of *subbridges*. A subbridge of a graph G with respect to components C_1, \dots, C_k is a maximal connected subgraph of G that does not become disconnected by removing all vertices of one component C_i ; see Figure 11.

Recall that we have to deal with bridges in two ways. First, each 2-component forms a SEFE instance with its $\textcircled{1}$ -bridges, while the union bridges partition these $\textcircled{1}$ -bridges (yielding union bridge constraints). Second, each union bridge yields a union bridge graph, which is a simpler instance one has to solve. In both cases, one can deal with subbridges instead of the whole bridges for the following reason. For the first case, we need the $\textcircled{1}$ -bridges only to create the corresponding dummy bridges. Thus, it suffices to know their attachment vertices. It is readily seen that each bridge B of C_i contains a

unique subbridge S incident to C_i and that the attachments of S at C_i are exactly the attachments of B at C_i . As this also holds for the union bridges, the union subbridges already define the correct grouping of the ①-subbridges. Concerning the second case, the SEFE instances that remain after exhaustively applying Lemma 9 are exactly the union subbridges together with a set of cycles, one for each incident 2-component. To conclude, it remains to show that each of the following three steps runs in linear time.

1. Compute for each 2-component the number of incident ①- and ②-subbridges, for each such subbridge its attachments, and the grouping of these subbridges into union subbridges.
2. Solve SEFE with union bridge constraints on instances with biconnected common graph.
3. Compute for each union subbridge a corresponding instance where each 2-component has been replaced by a suitable cycle.

For step 1 (Section 4.1), we contract every 2-component of G^{\cup} into a single vertex. The union subbridges are then basically the split components with respect to the resulting vertices. The same holds for the ①- and ②-subbridges (using G^{\circledone} and G^{\circledtwo} instead of G^{\cup}). For step 2 (Section 4.3), we modify the algorithm due to Angelini et al. [1]. Augmenting it such that it computes admissible embeddings in polynomial time is straightforward. Achieving linear running time is quite technical and, like the linear version of the original algorithm, requires some intricate data structures. For Step 3 (Section 4.2), computing the union subbridges is easy. To compute a suitable cycle for each incident 2-component C_i , one can make use of the fact that we already know an admissible embedding of C_i from step 2.

Theorem 5. *Given a SEFE instance, an equivalent set of instances of total linear size such that each 2-component of these instances is trivial can be computed in linear time.*

3.4 Special Bridges and Common-Face Constraints

In Section 3.3, we considered the case that C is a 2-component of the common graph G . We called the split components with respect to the vertices of C bridges (excluding edges in C). Clearly, this definition extends to the case where C is an arbitrary connected component. However, the decomposition into smaller instances does not extend to this more general case as for example Lemma 8 fails for non-biconnected graphs. Nonetheless, we are able to eliminate some special types of bridges in exchange for so-called common-face constraints. The reduction we describe in this section is thus in a sense weaker than the previous reductions as we reduce a given SEFE instance to a set of equivalent instances with common-face constraints. Many algorithmic approaches allow an easy integration of additional common-face constraints (see Section 5.3) and hence the more simple instances resulting from the reduction outweigh the disadvantages caused by the additional constraints (see Section 5.4).

Let $(G^{\circledcirc}, G^{\circledast})$ be an instance of SEFE with common graph G and let $\mathcal{F} \subseteq 2^{V(G)}$ be a family of sets of common vertices. A given SEFE satisfies the *common-face constraints* \mathcal{F} if and only if G has a face incident to all vertices in V' for every $V' \in \mathcal{F}$. The common-face constraints \mathcal{F} are *block-local* if for every $V' \in \mathcal{F}$ all vertices in V' belong to the same block of G .

Similarly, we say that a union bridge B is *block-local* if all attachment vertices of B belong to the same block of C . Let $B_1^{\circledcirc}, \dots, B_k^{\circledcirc}$ be the \circledcirc -bridges (for $i \in \{1, 2\}$) belonging to B . We say that B is *exclusively one-attached* if B_j^{\circledcirc} has only a single attachment vertex for $j = 1, \dots, k$.

Let B be a block-local union bridge of the common connected component C . Then the attachment vertices of B appear in the same order in every cycle of C (Lemma 8). Thus, we can define the union bridge graph G_B of B as in Section 3.3. Consider the SEFE instance $(H^{\circledcirc}, H^{\circledast})$ obtained from $(G^{\circledcirc}, G^{\circledast})$ by removing the union bridge B (the attachment vertices are not removed). It follows from Section 3.3 that $(G^{\circledcirc}, G^{\circledast})$ admits a SEFE if and only if the union bridge graph G_B admits a SEFE, and $(H^{\circledcirc}, H^{\circledast})$ admits a SEFE with an assignment of B to one of its faces f such (i) all attachment vertices of B are incident to f , and (ii) for $i \in \{1, 2\}$, no \circledcirc -bridge in B alternates with another \circledcirc -bridge in f .

If B is not only block-local but also exclusive one-attached, the latter requirement is trivially satisfied (a \circledcirc -bridge that has only a single attachment vertex cannot alternate). Thus, it remains to test whether G_B admits a SEFE and $(H^{\circledcirc}, H^{\circledast})$ admits a SEFE with block-local common-face constraints. We obtain the following theorem. The linear running time can be shown as in Section 4.

Theorem 6. *Given a SEFE instance, an equivalent set of instances with block-local common-face constraints of total linear size can be computed in linear time such that each instances satisfies the following property. No union bridge of a common connected component that is not a cycle is block-local and exclusively one-attached.*

4 Preprocessing 2-Components in Linear Time

As promised in the end of Section 3.3, we prove in this section that the decomposition of a SEFE instance into equivalent instances where every 2-component is a cycle can be done in liner time. Readers who want to skip this section can continue with Section 5 on page 40.

4.1 Computing the Sefe-Instances with Union Bridge Constraints

We first consider a slightly more general setting. Let $G = (V, E)$ be a graph and let C_1, \dots, C_k be disjoint connected subgraphs of G . We are interested in computing the number of bridges of each connected component C_i together with the attachments to C_i . We show that this can be done in $O(n + m)$ time (even if G is non-planar), where $n = |V|$ and $m = |E|$. Instead of computing directly the bridges and their attachments, our goal is rather to label each edge e that is incident to a vertex of some C_i but does not belong

to any of the C_i itself, by the bridge of C_i that contains e . Observe that, if each such incident edge has been labeled, the information about the number of C_i -bridges and their attachments can easily be extracted by scanning all incidences of vertices of C_i for $i = 1, \dots, k$. This scanning process can clearly be performed in total $O(n + m)$ time. In the following, we thus focus on computing this incidence labeling. Note that, since we are only interested in the attachments of bridges, it suffices to consider the corresponding subbridges as they have the same attachment sets.

Recall that a subbridge is a maximal connected subgraph of G for which none of the C_i is a separator. Note the high similarity of the definition of subbridges and the blocks of a graph, which are maximal connected subgraphs for which no single vertex is a separator. As with the blocks of a graph, it is readily seen that each edge of G that is not contained in one of the C_i is contained in exactly one subbridge of G . We exploit this similarity further and define the *component-subbridge tree* of G with respect to C_1, \dots, C_k as the graph that contains one vertex c_i for each component C_i and one vertex s_j for each subbridge S_j . Two vertices c_i and s_j are connected by an edge if and only if the subbridge S_j is incident to the component C_i . Note that, indeed, the component-subbridge tree is a tree. Once the component-subbridge tree has been computed, we can label each edge of $E \setminus (\bigcup_{i=1}^k C_i)$ with the subbridge containing it.

Lemma 11. *The component-subbridge tree of a graph G with respect to disjoint connected subgraphs C_1, \dots, C_k can be computed in linear time.*

Proof. First, contract each component C_i to a single vertex c_i ; call the resulting graph G' . Note that, in G' , the subbridges are exactly the maximal connected subgraphs for which none of the vertices c_i is a separator. We compute the component-subbridge tree T in three steps. First, compute the block-cutvertex tree of G' . Second, for each cutvertex v that does not correspond to one of the c_i , remove v and merge its incident blocks into the same subbridge. Finally, create for each component C_i that has only one bridge a corresponding vertex c_i and attach it as a leaf to the unique subbridge incident to C_i . Clearly, each of the steps can be performed in $O(n + m)$ time. \square

As argued above, Lemma 11 can be used to label in linear time the incident edges of the components C_1, \dots, C_k by their corresponding bridges. For step 1 of our reduction, we take C_1, \dots, C_k as the 2-components of a SEFE instance (G^\oplus, G^\otimes) . We then use the above approach to label the attachment incidences of the ①-, ②-, and the union bridges of C_1, \dots, C_k . From this we can create the dummy bridges for each 2-component C_i together with the union bridge constraints in time linear in the sum of degrees of vertices in C_i . By the arguments for Lemma 10, the resulting instance admits a SEFE if and only if C_i has an admissible embedding. Since the C_i are disjoint, it follows that the construction of all instances can be done in linear time. This finishes step 1.

4.2 Constructing the Subbridge Instances

Let us assume that each 2-component has an admissible embedding, which is found using the linear-time algorithm described in Section 4.3. Otherwise a SEFE of the original

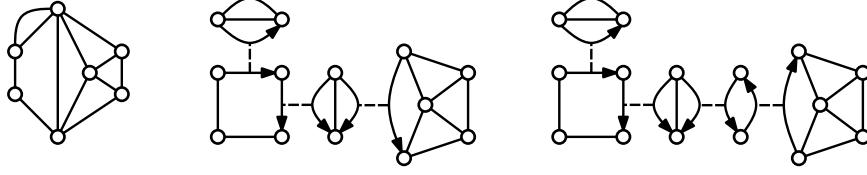


Figure 12. A graph (left) together with its SPQR-tree (middle). The Q-nodes are omitted to improve readability. The augmented SPQR-tree (right) contains an additional S-node whose skeleton is a cycle of length 2.

graph does not exist. In the final step of our reduction, we substitute, in each subbridge, the incident 2-components by a cycle. This results in a set of SEFE instances—one for each subbridge—that all admit a solution if and only if the original instance admits a solution. They can hence be handled completely independently. To efficiently extract all instances, we process the 2-components independently and replace each one by a cycle in their incident subbridges. The time to process a single 2-component with all its incident subbridges is linear in the size of the 2-component plus the number of attachments of these subbridges in the respective 2-component. It then immediately follows that processing all 2-components in this way takes linear time.

Consider a fixed 2-component C with an admissible embedding as computed in step 2 of the reduction. Consider a fixed face f together with the subbridges that are embedded in that face. For each bridge B embedded in f , we construct a list of attachments A_B . We traverse the facial cycle of f . At each vertex, we check the edges embedded inside this face and append the vertex to the list of each subbridge for which it is an attachment. Afterwards, we traverse for each subbridge its list of attachments and replace C by a cycle that visits the attachments in the order of the attachment list. The time is clearly proportional to the size of f and the attachments of the subbridges embedded in f . Hence processing all faces of all components in this way takes linear time and yields the claimed result. This implements step 3 in linear time.

4.3 Simultaneous Embedding with Union Bridge Constraints

In this section, we show how to solve SEFE with union bridge constraints in linear time if the common graph is biconnected. Our algorithm is based on the algorithm by Angelini et al. [1]. Note that our extension to this algorithm is twofold. We allow bridges with an arbitrary number of attachment vertices. The original algorithm allows only two attachments per ①-bridge (i.e., each bridge is a single exclusive edge). Moreover, we have to deal with union bridge constraints. To avoid some special cases and simplify the description, we sometimes deviate from the notation used by Angelini et al. Our focus lies on a linear-time implementation, the correctness of our approach directly follows from the correctness of the algorithm by Angelini et al.

Let G be the biconnected common graph. We consider the (unrooted) *augmented SPQR-tree* \mathcal{T} of G , which is defined as follows; see Figure 12. Let \mathcal{T}' be the SPQR-tree of G and let μ_1 and μ_2 be two adjacent nodes in \mathcal{T}' such that each of them is a P- or

an R-node. We basically insert a new S-node between μ_1 and μ_2 whose skeleton is a cycle of length 2 (i.e., a pair of parallel virtual edges). More precisely, let $\varepsilon_1 = \{s, t\}$ and $\varepsilon_2 = \{s, t\}$ be the virtual edges in μ_1 and μ_2 , respectively, that correspond to each other. We subdivide the edge μ_1, μ_2 in \mathcal{T}' ; let μ be the new subdivision vertex. The skeleton $\text{skel}(\mu)$ contains the vertices s and t with two virtual edges ε'_1 and ε'_2 between them corresponding to ε_1 in $\text{skel}(\mu_1)$ and ε_2 in $\text{skel}(\mu_2)$, respectively. Applying this augmentation for every pair of adjacent nodes that are P- or R-nodes gives the augmented SPQR-tree \mathcal{T} . Note that P- and R-nodes in \mathcal{T} have only S- and Q-nodes as neighbors. Moreover, no two S-nodes are adjacent.

Consider a bridge B and let μ be a node of \mathcal{T} . A virtual edge ε in $\text{skel}(\mu)$ is an *attachment* of B if its expansion graph contains an attachment vertex of B . We say that B is *important* for μ if it has at least two distinct attachments among the vertices and virtual edges of $\text{skel}(\mu)$ that are not two adjacent vertices in $\text{skel}(\mu)$. It is clearly necessary, that $\text{skel}(\mu)$ admits an embedding such that for every union bridge B the attachments in $\text{skel}(\mu)$ are incident to a common face. An embedding having this property is called *compatible*. This leads to the following first step of the algorithm.

Step 1: Compatible embeddings. For every P- and R-node μ , compute the important union bridges with their attachments. If μ is an R-node, check whether the unique (up to flip) embedding $\text{skel}(\mu)$ is compatible. If μ is a P-node check whether it admits a compatible embedding and fix such an embedding (up to flip).

If Step 1 fails, the instance does not admit a SEFE. Note that the skeleton of a P-node might admit several compatible embeddings. However, fixing an arbitrary compatible embedding up to flip does not make a solvable SEFE instance unsolvable [1]. Thus, after Step 1, we can assume that the embedding of every skeleton is fixed and it remains to decide for each skeleton whether its embedding should be flipped or not. We call the embedding fixed for a skeleton its *reference embedding*.

For every P- and R-node μ let x_μ be a binary decision variable with the following interpretation. The skeleton $\text{skel}(\mu)$ is embedded according to its reference embedding and according to the flipped reference embedding if $x_\mu = 0$ and $x_\mu = 1$, respectively. By considering the S-nodes of the augmented SPQR-tree, one can derive necessary conditions for these variables that form an instance of 2-SAT (actually, we only get equations and inequalities, which is a special case of 2-SAT).

Let μ be an S-node. We assume the edges in $\text{skel}(\mu)$ to be oriented such that $\text{skel}(\mu)$ is a directed cycle. Thus, we can use the terms left face and right face to distinguish the faces of $\text{skel}(\mu)$. Let B be a bridge that is important for μ . We can either embed B into the left or into the right face of $\text{skel}(\mu)$. We define the binary variable x_B^μ with the interpretation that B is embedded into the right face and into the left face of $\text{skel}(\mu)$ if $x_B^\mu = 0$ and $x_B^\mu = 1$, respectively.

Assume that the virtual edge $\varepsilon = st$ (oriented from s to t) in $\text{skel}(\mu)$ is an attachment of B , i.e., an attachment vertex $v \notin \{s, t\}$ of B lies in the expansion graph of ε . Let μ' be the neighbor of μ corresponding to ε and let ε' be the twin of ε in $\text{skel}(\mu')$ (also oriented from s to t); see Figure 13. Clearly, B is also important for μ' as ε' contains an

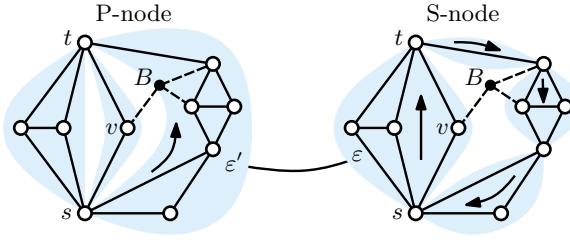


Figure 13. An S-node with a bridge B having ε as right-sided attachment. The virtual edges are illustrated as gray regions with their expansion graph inside.

attachment vertex of B while the attachment vertex v is not contained in ε' . In Step 1 we ensured that $\text{skel}(\mu')$ is embedded such that v (or the virtual edge containing v) shares a face with ε' . If this face lies to the left of ε' in $\text{skel}(\mu')$, we say that v is an attachment *on the right side* of ε in $\text{skel}(\mu)$ (as in the example in Figure 13). Otherwise, if this face lies to the right of ε' , we say that v is an attachment *on the left side* of ε . For a bridge B with attachment vertex v we also say that the attachment ε is *right-sided* and *left-sided* if v lies on the right and left side of ε , respectively.

Assume B has an attachment on the right side of the virtual edge ε in the skeleton $\text{skel}(\mu)$ (as in Figure 13). Assume further that the skeleton $\text{skel}(\mu')$ of the corresponding neighbor is not flipped, i.e., $x_{\mu'} = 0$. Then B must be embedded into the face to the right of the cycle $\text{skel}(\mu)$, i.e., $x_B^\mu = 0$. Conversely, if the embedding of $\text{skel}(\mu')$ is flipped, i.e., $x_{\mu'} = 1$, then B must lie in the left face of $\text{skel}(\mu)$, i.e., $x_B^\mu = 1$. This necessary condition is equivalent to the equation $x_{\mu'} = x_B^\mu$. Similarly, if B has an attachment on the left side of ε , we obtain the inequality $x_{\mu'} \neq x_B^\mu$. We call the resulting set of equations and inequalities the *consistency constraints* of the bridge B in μ . This leads to the second step of the algorithm.

Step 2: Consistency constraints. For every S-node μ compute the important ①-bridges (for $i \in \{1, 2\}$) and union bridges together with their attachments. For attachments in virtual edges also compute whether they are left- or right-sided. Then add the consistency constraints of these bridges in μ to a global 2-SAT formula.

The consistency constraints are necessary but not sufficient as they do not ensure that no two alternating bridges of the same type are embedded into the same face. Consider an S-node μ with two important bridges B and B' . Assume these two bridges alternate (i.e., they have alternating attachments in the cycle $\text{skel}(\mu)$). Embedding B and B' on the same side of $\text{skel}(\mu)$ yields a crossing between an edge in B and an edge in B' . Thus, if B and B' are both ①- or both ②-bridges, then they must be embedded to different sides of $\text{skel}(\mu)$. In this case, we obtain the inequality $x_B^\mu \neq x_{B'}^\mu$. This inequality is called *planarity constraint*.

Step 3: Planarity constraints. For every S-node μ compute the pairs of important ①-bridges that alternate in $\text{skel}(\mu)$. Do the same for ②-bridges. For each such pair add the planarity constraint to the global 2-SAT formula.

Finally, we have to embed ①-bridges belonging to the same union bridge into the same face. Let B be an ①-bridge and let B' be the union bridge it belongs to. The *union-bridge constraint* of B in μ is the equations $x_B^\mu = x_{B'}^\mu$.

Step 4: Union-bridge constraints. For every S-node μ , add the union-bridge constraint of each important ①-bridge to the global 2-SAT formula.

After Steps 2–4, the global 2-SAT formula is solved in linear time [2]. The solution determines for every P- and every R-node μ , whether the reference embedding of $\text{skel}(\mu)$ should be flipped or not, which completely fixes the embedding of the common graph G . Of course, there might be different solutions of the 2-SAT formula, yielding different embeddings. However, if one of these solutions yields a SEFE, then any of the solutions does [1]. Thus, one can simply take one solution and check whether it yields a SEFE (with union bridge constraints) or not.

Step 5: Final step. Test whether the given instance admits a SEFE with union bridge constraints assuming that the embedding of the common graph is fixed.

It remains to implement Steps 1–5 in linear time, which is done in the following. We first note that there are too many important bridges to be able to compute them in linear time (as required in Step 1 and Step 2). However, similar to Angelini et al. [1], we can show that many important bridges can be omitted without loosing the correctness of the algorithm, which leads to a linear-time implementation.

4.3.1 Too Many Bridges are Important

We start with the observation that computing all important union bridges for every P- and R-node of the SPQR-tree is actually a bad idea, as there may be $\Omega(n)$ bridges each being important in $\Omega(n)$ nodes. Thus, explicitly computing all of them would require $\Omega(n^2)$ time. Consider the graph G in Figure 14 whose SPQR-tree \mathcal{T} (without Q-nodes) is a path. Let μ be one of the P-nodes (note that $\text{skel}(\mu)$ has two virtual edges and one normal edge) and let B be one of the bridges shown in Figure 14. Clearly, the expansion graphs of both virtual edges of $\text{skel}(\mu)$ contain at least one attachment vertex of B . Thus, B is important for μ and B has the two virtual edges of $\text{skel}(\mu)$ as attachments. As this may hold for a linear number of bridges, we get the above observation.

To resolve this issue, note that from the perspective of μ (in the above example), all bridges look the same in the sense that they have the same set of attachments. Intuitively, they thus lead to similar constraints and it seems to suffices to know only one of these bridges. In the following, we first show that omitting some of the bridges is indeed safe in the sense that the algorithm remains correct. Then we show how to compute the remaining bridges efficiently.

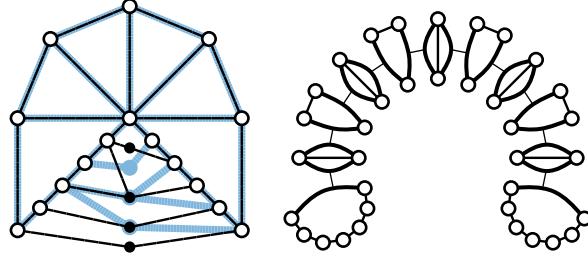


Figure 14. A graph with many bridges (left) each of which being important for every node of the SPQR-tree (right).

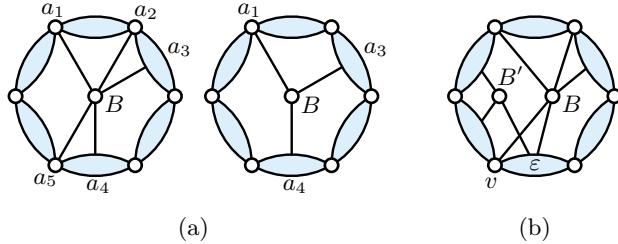


Figure 15. (a) The bridge B has five attachments a_1, \dots, a_5 . Attachments a_2 and a_5 are superfluous due to a_3 and a_4 , respectively. (b) The bridges B and B' alternate only due to the superfluous attachment v of B . However, the consistency constraints already synchronize B and B' as they have ε as common attachment.

4.3.2 Omitting Some Important Bridges

In this section we show how to change the algorithm described above (Steps 1–5) slightly without changing its correctness. We say it is *safe* to do something if doing it does preserve the correctness of the algorithm. In particular, we show that it is safe to omit some of the important bridges. In the subsequent sections we then show that the remaining important bridges can be computed efficiently leading to an efficient implementation of all five steps.

Let G be the common graph and let \mathcal{T} be its SPQR-tree. Let μ be an inner node of \mathcal{T} and let B be a bridge that is important for μ with attachments a_1, \dots, a_ℓ in $\text{skel}(\mu)$ (recall that each of the a_i is either a vertex or a virtual edge of $\text{skel}(\mu)$). We call an attachment a_i (with $1 \leq i \leq \ell$) *superfluous*, if a_i is a vertex in $\text{skel}(\mu)$ such that B has another attachment a_j that is a virtual edge incident to the vertex a_i ; see Figure 15a. The following lemma shows that the term ‘‘superfluous’’ is justified.

Lemma 12. *Omitting superfluous attachments is safe.*

Proof. There are two situations in which missing superfluous attachments might play a role. First, when we check a P- or R-node skeleton for a compatible embedding (Step 1). Second, when we add the planarity constraints for alternating ①-bridges (Step 3). Let v be the superfluous attachment of B and let ε be a virtual edge incident to v that is also an attachment of B in $\text{skel}(\mu)$. Concerning compatible embeddings, we have to make sure

that $\text{skel}(\mu)$ admits an embedding where all attachments of B are incident to a common face. Clearly, v is incident to both faces the virtual edge ε is incident to. Thus, omitting the attachment v does not change anything.

Concerning the planarity constraints, we have to consider the case that B alternates with another bridge B' only due to the attachment v in B . This can only happen if B' has ε as attachment; see Figure 15b. However, then the consistency constraints (Step 2) either forces B and B' into different faces (if their attachment in ε lies on different sides of ε) and everything is fine, or they force B and B' to lie in the same face. In the latter case, the instance is clearly not solvable, which will be found out in Step 5. \square

In the following we always omit superfluous attachments even when we do not mention it explicitly. Note that this retroactively changes the definition of important bridges slightly, i.e., a bridge B is important for a node μ if B has at least two (non-superfluous) attachments in $\text{skel}(\mu)$.

To show that we can omit sufficiently many important bridges to get a linear running time, we have to root the SPQR-tree \mathcal{T} . More precisely, we choose an arbitrary Q-node as the root of \mathcal{T} .

We categorize the important bridges in different types of bridges depending on their attachments. To this end, we first define different types of attachments. Let μ be a node of the (rooted) SPQR-tree and let B be an important bridge of μ with attachments a_1, \dots, a_ℓ . Recall that an attachment is either a vertex of a virtual edge of $\text{skel}(\mu)$. If a_i is a pole of $\text{skel}(\mu)$, we call it a *pole attachment*. If a_i is the parent edge of $\text{skel}(\mu)$, we call it *parent attachment*. All other attachments are called *child attachments*.

We say that the important bridge B is a *regular bridge* of μ if B has at least two child attachments. If B has only a single child attachment, it has either a parent attachment or one or two pole attachments (note that pole attachments are superfluous in the presence of a parent attachment). We call B *parent bridge* and *pole bridge* in the former and latter case, respectively. Note that B must have at least one child attachment as it otherwise cannot be important.

As shown before, we cannot hope to compute all important bridges efficiently as there may be too many of them. Thus, we show in the following that omitting some important bridges is safe.

Lemma 13. *Let B be a parent bridge whose child attachment a is a virtual edge. Omitting another parent bridge or a pole bridge with child attachment a (while keeping B) is safe.*

Proof. Let B' be an other important bridge of μ and let a be the child attachment of B and B' . Let ε be the parent edge of $\text{skel}(\mu)$. In Step 1, the bridge B forces the embedding of $\text{skel}(\mu)$ to have a and ε on a common face. If B' is a parent bridge, it requires the same (and thus no additional) condition for the embedding of $\text{skel}(\mu)$. If B' is a pole bridge, it requires one of the poles (or both of them) to be incident to a common face with a . However, this is clearly the case if the parent edge ε (which is incident to both poles) shares a face with a . Hence, omitting B' does not change anything in Step 1.

We cannot argue separately for Steps 2–4 as they all contribute to the same global 2-SAT formula. We call the 2-SAT formula we get when omitting B' NEW. The formula

we get when not omitting B' is called ORIGINAL. The straightforward way to prove that omitting B' is safe would be to show that the new and the original 2-SAT formulae are equivalent in the sense that they have the same solutions. However, this is not true as omitting B' can make the 2-SAT formula solvable whereas it was unsolvable before.

Assume the original 2-SAT formula is not solvable. Then the given instance does not admit a SEFE with union bridge constraints. In this case, Step 5 will never succeed no matter what we do in the steps before. Thus, we only have to argue for the case that the original 2-SAT formula is solvable. In this case, the only thing that can go wrong is that the new 2-SAT formula admits a solution that is not a solution in the original formula. This solution could then result in an embedding of the common graph G that does not admit a SEFE, whereas all solutions of the original 2-SAT formula lead to a SEFE.

To get a handle on this, we consider the conflict graph of a 2-SAT formula. First note that the formula we obtain is special in the sense that each constraint is either an equation or an inequality. We define the conflict graph to have a vertex for each variable and an edge connecting two vertices if there is an (in-)equality constraint between the corresponding variables. Recall that the 2-SAT formula has a variable x_μ for every P- or R-node μ indicating whether the embedding of $\text{skel}(\mu)$ has to be flipped or not. In the following we show that two such variables x_μ and x_η that are in the same connected component of the conflict graph of the original 2-SAT formula are also in the same connected component with respect to the new formula. This shows that every embedding resulting from a solution of the new formula also yields a solution of the original formula (assuming that the original formula has a solution at all).

First note that omitting an important bridge B' in an S-node μ has the effect that the vertex corresponding to the variable $x_{B'}^\mu$, is deleted in the conflict graph (for convenience we also denote the vertex by $x_{B'}^\mu$). Recall that the variable $x_{B'}^\mu$ represents the decision of whether B' is embedded into the left or the right face of $\text{skel}(\mu)$. We show that $x_{B'}^\mu$ is dominated by x_B^μ in the sense that every neighbor of $x_{B'}^\mu$ is also connected to x_B^μ by a path not containing $x_{B'}^\mu$. Thus, removing $x_{B'}^\mu$ does not change the connected components of the conflict graph.

Consider the consistency constraints of B' (Step 2). Let a be an attachment of B' . If a is a vertex in $\text{skel}(\mu)$, we do not get a consistency constraint there. If a is a virtual edge, it corresponds to a neighbor η of μ and we get the constraint $x_\eta = x_{B'}^\mu$ or $x_\eta \neq x_{B'}^\mu$. But then B has also a as attachment and thus we also have one of the two constraints $x_\eta = x_B^\mu$ or $x_\eta \neq x_B^\mu$.

For the planarity constraints (Step 3), first assume that B' is a parent bridge. Then B alternates with another bridge B'' if and only if B' alternates with B'' . Thus, if we have the constraint $x_{B'} \neq x_{B''}$ we also have the constraint $x_B \neq x_{B''}$. If B' is a pole bridge, there might be a bridge B'' alternating with B' (yielding the constraint $x_{B'} \neq x_{B''}$), whereas B and B'' do not alternate. However, this can only happen if B'' has the parent edge ε of $\text{skel}(\mu)$ as attachment. Let η be the neighbor of μ corresponding to ε (i.e., the parent of μ). From Step 2 we then have the consistency constraint between x_η and $x_{B''}^\mu$ and also one between x_η and x_B^μ (B is a parent bridge). Thus, the conflict graph includes the path $x_B^\mu x_\eta x_{B''}^\mu$, which still exists after removing B' .

Finally, if B'' is the union-bridge including B' , we loose the union-bridge constraint

(Step 4) $x_{B'}^\mu = x_{B''}^\mu$ by omitting B' . Note that the attachments of B' are a subset of the attachments of B'' . Thus, B'' has the child attachment a (which is a virtual edge). Let η be the neighbor of μ corresponding to a . Then we have a consistency constraint connecting $x_{B''}^\mu$ with x_η . Moreover, x_η is also connected to X_B^μ as a is an attachment of B . This yields a connection from x_B^μ to $x_{B''}^\mu$. \square

This lemma gives rise to the following definition. Let B' be a pole bridge of μ with child attachment a that is a virtual edge in $\text{skel}(\mu)$. Let further B be a parent bridge of μ with child attachment a . We say that B' is *dominated* by B . If B' is a parent bridge of μ with child attachment a instead, we say that B and B' are *equivalent*. Note that being equivalent is clearly an equivalence relation. Lemma 13 shows that we can omit dominated pole bridges and all but one parent bridge for each equivalence class.

4.3.3 Computing the Remaining Important Bridges

In this section we show that all bridges that are not omitted due to Lemma 13 can be computed in linear time. Actually, we compute slightly more information, which we need in some intermediate steps. We for example never omit a parent bridge B in μ if B is regular in the parent of μ . We call such a bridge *semi-regular* in μ .

Moreover, consider an important bridge B of μ and let a be an attachment of B in $\text{skel}(\mu)$ that is a virtual edge. A vertex v of G that is not incident to a in $\text{skel}(\mu)$ is a *representative* of the attachment a if v is an attachment vertex of B and included in the expansion graph of a . If the attachment a is a vertex of $\text{skel}(\mu)$, we say that a is its own representative. Note that every attachment of B has at least one representative vertex. When we compute the important bridges of a node μ , we also compute their attachments in μ together with at least one representative for each attachment. Before we actually compute bridges, we provide some general tools.

Lemma 14. *After linear preprocessing time, the pole and parent attachments of a bridge in a given node together with a representative of each attachment can be computed in constant time.*

Proof. We first do some preprocessing. We define the *SPQR-vertex-tree* \mathcal{T} to be the tree obtained from the SPQR-tree by removing its Q-nodes and attaching every vertex v of G as a leaf to the highest node that contains v in its skeleton. Note that this is the unique node that contains v in its skeleton but not as pole. For a non-pole vertex v in a skeleton $\text{skel}(\mu)$, we say that the leaf v (which is a child of μ) corresponds to v . Note that this makes sure that every child attachment in $\text{skel}(\mu)$ corresponds to a child of μ in \mathcal{T} .

Assume the vertices of G (i.e., the leaves of \mathcal{T}') to be numbered according to a DFS-ordering (which we get in $O(n)$ time). We start by sorting the attachment vertices of each bridge according to this DFS-ordering. For all bridges B_1, \dots, B_k together, this can be done in time $O\left(n + \sum_{i=1}^k |B_i|\right)$ as follows. To simplify the notation, we identify every bridge B_i with its set of attachment vertices. First, sort all pairs (v, B_i) with $v \in B_i$ (i.e., basically the disjoint union of the B_i) using bucket sort with buckets $1, \dots, n$ (for each

vertex v , one bucket containing all pairs (v, B_i)). It then suffices to iterate over all these sorted pairs to extract the sets B_1, \dots, B_k sorted according to this DFS-ordering.

For a node of \mathcal{T} , the leaves that are descendants of this node appear consecutively in the DFS-ordering. By going bottom-up in \mathcal{T} , we can compute for every node μ the leaf with the smallest and the leaf with the largest number among the descendants of μ . Thus, given a vertex v of G , we can decide in constant time whether v is a descendent of μ .

Now we answer the queries. Let B be a bridge that is important in a node μ . If the vertex with the smallest and the vertex with the largest number in B (which we can find in constant time as B is sorted) are both descendants of μ , all attachment vertices of B are descendants of μ . In this case B has neither a pole nor a parent attachment in μ . Otherwise, by looking at the first two and the last two elements in B we can distinguish the following two cases. (i) There is an attachment vertex v that is not a descendent of μ and not a pole of $\text{skel}(\mu)$. Then the parent edge of $\text{skel}(\mu)$ is an attachment of B in $\text{skel}(\mu)$ and v is a representative for this attachment. If B also has pole attachments in $\text{skel}(\mu)$, they are superfluous and we can ignore them by Lemma 12. (ii) There is no such vertex v . Then all attachment vertices in B are descendants of μ except for maybe the two poles. In this case, the poles of $\text{skel}(\mu)$ that are attachments of B are among the first or last two vertices in B and thus we find all pole attachments of B in $\text{skel}(\mu)$. This concludes the proof. \square

Lemma 15. *The regular bridges and their attachments together with a representative for each attachment can be computed in linear time.*

Proof. Let \mathcal{T} be the SPQR-vertex-tree of G and assume again that the vertices in every bridge are sorted according to a DFS-order of the leaves in \mathcal{T} . We first show how to compute all nodes in which a given bridge B_i is regular in $O(|B_i|)$ time. Note that this implicitly shows that B_i has only $O(|B_i|)$ regular vertices.

The *lowest common ancestor (LCA)* of two vertices u and v is the highest node on the path between u and v . We denote it by $\text{LCA}(u, v)$. Clearly, the vertices u and v are descendants of different children of $\text{LCA}(u, v)$. Thus, if B_i has the attachment vertices u and v , then B_i has two different child attachments in $\text{LCA}(u, v)$ with representatives u and v . Thus, $\text{LCA}(u, v)$ is active. Conversely, if B_i is regular in a node μ , it has two different child attachments. Let u and v be representatives of these two attachments. Clearly, u and v are descendants of different children of μ and thus $\mu = \text{LCA}(u, v)$ holds. Hence, the nodes in which B_i is regular are exactly the LCAs of pairs of attachment vertices in B_i .

To see that it is not necessary to consider all pairs of vertices in B_i , let u , v , and w with $u < v < w$ (according to the DFS-ordering) be contained in B_i . Then $\text{LCA}(u, w) = \text{LCA}(u, v)$ or $\text{LCA}(u, w) = \text{LCA}(v, w)$ holds for the following reason. If $\mu = \text{LCA}(u, v) = \text{LCA}(v, w)$, then u , v , and w are descendants of three different children of μ . Thus, $\text{LCA}(u, w) = \mu$ also holds. Otherwise, assume $\mu = \text{LCA}(u, v)$ is a descendant of $\eta = \text{LCA}(v, w)$. Thus, from the perspective of η , the vertices u and v and the node μ are descendants of the same child whereas w is the descendant of a different child. Thus $\text{LCA}(u, w) = \text{LCA}(v, w) = \eta$. Hence, to compute all nodes in which B_i is regular, it

suffices to compute the LCA for pairs of attachment vertices in B_i that are consecutive (with respect to the ordering we computed before). As the LCA can be computed in constant time after $O(n)$ preprocessing time [14], this gives us all regular bridges of all nodes of the SPQR-tree in overall $O(n + \sum_{i=1}^k |B_i|)$ time.

Given a node μ and a regular bridge B of μ , we are still lacking the attachments of B in $\text{skel}(\mu)$. We can get the parent and pole attachments using Lemma 14. Consider a child attachment a of B in $\text{skel}(\mu)$. Note that the attachment vertices of B that are descendants of a are consecutive with respect to the DFS-order. When choosing the first or last of these vertices, we get an attachment vertex v in B with the following properties: v is a descendent of the child of μ that corresponds to a and either $\mu = \text{LCA}(u, v)$ for the predecessor u of v in B (ordered according to the DFS-order) or $\mu = \text{LCA}(v, w)$ for the successor w of v . Thus, we actually already know a representative for every child attachment of B in $\text{skel}(\mu)$. To find for an attachment vertex v (the representative) the corresponding attachment in $\text{skel}(\mu)$, we need to find the child of μ that has the leaf v as a descendant. This can be done for all active bridges simultaneously by processing \mathcal{T} bottom-up while maintaining a union-find data structure. As the sequence of union operations is known in advance, each union and find operation takes amortized $O(1)$ time [10]. Thus, it takes $O\left(n + \sum_{i=1}^k |B_i|\right)$ time in total. \square

Lemma 16. *The semi-regular bridges and their attachments together with a representative for each attachment can be computed in linear time.*

Proof. By Lemma 15, we know for every node μ all the bridges that are regular in μ . We can assume that the regular bridges of a μ are stored as a list that is sorted according to an arbitrarily chosen order of the bridges (we just have to process the bridges in this order in the proof of Lemma 15). Moreover, we know all the attachments of B in $\text{skel}(\mu)$ together with a representative for each attachment.

Let a be a child attachment of B in $\text{skel}(\mu)$ with representative v . Let η be the child of μ corresponding to a . If η is a leaf, it actually must be the vertex v and there is nothing to do. Otherwise, v is a descendent of η and thus B has a child attachment in η . Moreover, it also has a parent attachment in η as it otherwise would not be regular in the parent μ of η . Thus, B is either regular or a parent bridge (and hence semi-regular) in η .

By processing all regular bridges of μ like that, we can build for η (and all other children of μ) a list of bridges that contains all semi-regular bridges and additionally some regular bridges. Note that building up this list takes constant time for each attachment of regular bridges. As we computed those attachments in linear time, there cannot be more than that many attachments.

To get rid of the bridges that are actually regular and not semi-regular, note that we can assume the list of semi-regular and regular bridges to be sorted (according to the arbitrary order of bridges we chose before). Thus, we can simply process this list and the list of regular bridges of η (which we computed using Lemma 15) simultaneously and throw out all those that appear in both lists. This leaves us with a list of semi-regular bridges for every node. In addition to that, we know an attachment vertex for every semi-regular bridge that is a representative of the child-attachment of that bridge. Thus, we also get

the actual attachments in $\text{skel}(\eta)$ in linear time using one bottom-up traversal as in the proof of Lemma 15. Moreover, we get the parent attachments using Lemma 14. \square

Lemma 17. *The parent and pole bridges and their attachments together with a representative for each attachment can be computed in linear time when omitting dominated pole bridges and all but one parent bridge of each equivalent classes.*

Proof. Let \mathcal{T} be the SPQR-vertex-tree of the common graph G .

We now process \mathcal{T} bottom-up computing a list of bridges for each node μ . This list will contain all bridges that potentially are parent or pole bridges in μ and we denote it by $\text{pot}(\mu)$. More precisely, if a bridge B has an attachment vertex that is a descendent of μ and other attachment vertices that are not descendants of μ , then B is contained in $\text{pot}(\mu)$. As we do not have the time to tidy up properly, $\text{pot}(\mu)$ can also contain bridges whose attachment vertices are all descendants of μ .

We start with the leaves of \mathcal{T} , which are vertices of G . Let v be such a leaf. Then $\text{pot}(v)$ is the list of all bridges having v as attachment. By initializing $\text{pot}(v)$ with the empty list and then processing all bridges once, we can compute $\text{pot}(v)$ for all vertices in $O\left(n + \sum_{i=1}^k |B_i|\right)$ time. Now consider an inner node μ . We basically obtain the list $\text{pot}(\mu)$ by concatenating the lists of all children. But before concatenation we (partially) process these lists separately.

While processing μ (and the lists computed for the children of η) we want to answer for a given bridge B the following queries in constant time. First, is B regular in μ ? Second, have we seen B already while processing μ ? This can be done using timestamps. Assume we have one global array with an entry for each bridge. Before processing μ we increment a global timestamp and write this timestamp into the fields of the array corresponding to bridges that are regular in μ . While processing μ , we can then check in constant time whether the current timestamp is set for a given bridge B and thus whether B is regular. Setting up this array takes time linear in the number of regular bridges of μ and thus we have an overall linear overhead. We can handle the second query in constant time, analogously.

Now let μ be the node we currently process, let η be a child of μ and assume that $\text{pot}(\eta)$ is already computed. We process the list $\text{pot}(\eta)$; let B be the current bridge. By Lemma 14 we can check in constant time whether B has pole or parent attachments in μ . If not, all attachment vertices of B are descendants of μ and we remove μ from the list $\text{pot}(\eta)$ as B cannot be a parent or pole in μ or in any ancestor of μ .

Otherwise, B has a parent or a pole attachment in μ . We first check (in constant time) whether B is regular in μ . Assume it is regular and we see B the first time since processing μ (which we can also check in constant time as mentioned above). Then we simply skip B and continue processing $\text{pot}(\eta)$. If B is regular in μ and B occurred before while processing μ , it would be contained twice in the list $\text{pot}(\mu)$ when concatenating the lists of all children of μ . Thus we can remove this occurrence of B from $\text{pot}(\eta)$. Afterwards, we continue processing the remaining bridges in $\text{pot}(\eta)$.

Now assume that B is not regular. Then B is either a parent or a pole bridge. If B is a parent bridge, we store it as a parent bridge of μ . Note that we also know the attachments

of B in $\text{skel}(\mu)$ (the parent edge and the attachment corresponding to the child η) and a representative for each of these attachments. Afterwards, we stop processing $\text{pot}(\eta)$ and continue with another child of μ . By stopping after processing B , we might miss a bridge B' in $\text{pot}(\eta)$ that is also a pole or parent bridge of μ . However, the child attachment of B' would be the attachment corresponding to η and thus B' is in the same equivalence class as B (if B' is a parent bridge) or B' is dominated by B (if B' is a pole bridge). In both cases we can omit B' .

Finally, consider the case that B is a pole bridge in μ . We save B together with its attachments in $\text{skel}(\mu)$ (and their representatives) as pole bridge of μ . Then we remove B from $\text{pot}(\eta)$ and continue processing $\text{pot}(\eta)$. Removing B from $\text{pot}(\eta)$ has the effect that B does not occur when processing an ancestors of μ . Thus, we have to show that B is not a pole or parent bridge in one of these ancestors. Consider an ancestor τ of μ . Then either all attachment vertices of B are descendants of τ and B is neither pole nor parent bridge for τ . Otherwise, the only attachment vertex of B that is not a descendent of τ is s (without loss of generality). However, the virtual edge in $\text{skel}(\tau)$ containing all attachment vertices of B is then incident to s and thus the attachment s is superfluous. Hence, we can omit the bridge B in $\text{skel}(\tau)$ by Lemma 12.

It remains to show that the above procedure runs in linear time. First note that we add elements to lists $\text{pot}(\cdot)$ only in the leaves of \mathcal{T} . As we add each bridge B to exactly $|B|$ such bridges, the total size of these lists is linear. Assume we are processing a list $\text{pot}(\eta)$ and let B be a bridge we delete after processing it. Then we can ignore the (constant) running time for processing B , as we have only linearly many such deletion operations. Otherwise, B is a regular bridge that we see the first time or it is a parent bridge. The former case happens only as many times as there are regular bridges of μ (which is overall linear). The latter case happens at most once for each child of μ as we stop processing $\text{pot}(\eta)$ afterwards. Hence, the overall running time is linear. \square

4.3.4 Linear Time Implementation of Steps 1–5

Lemma 18. *Step 1 can be performed in linear time.*

Proof. Recall that Step 1 consist of computing the important union bridges for P- and R-nodes. For every P- and R-node μ we then have to test whether $\text{skel}(\mu)$ admits a compatible embedding, i.e., whether $\text{skel}(\mu)$ can be embedded such that the attachments of each important bridge of μ share a face.

For each node μ , we compute the regular and semi-regular bridges using Lemma 15 and Lemma 16, respectively. We moreover compute some of the pole and parent bridges using Lemma 14. In this way we of course miss some important bridges but we know by Lemma 13 that it is safe to do so. Thus, we can focus on computing compatible embeddings.

Let μ be a node of the SPQR-tree and assume that μ is a P-node. Each of the two vertices s and t in $\text{skel}(\mu)$ is incident to every face. If an important bridge B has s or t as attachment, this attachment does not constrain the embedding of $\text{skel}(\mu)$ (it shares a face with all other attachments of B if and only if all other attachments share a face).

Thus, we can assume that only the virtual edges of $\text{skel}(\mu)$ are attachments. If B has three (or more) attachments in $\text{skel}(\mu)$, it is impossible to find a compatible embedding as every face of $\text{skel}(\mu)$ is incident to only two virtual edges. It remains to deal with the case where every bridge has two virtual edges as attachment. We build the conflict graph with one vertex $v(\varepsilon)$ for every virtual edge ε and an edge between two such vertices $v(\varepsilon_1)$ and $v(\varepsilon_2)$ if and only if there is a bridge with attachments ε_1 and ε_2 . It is not hard to see that $\text{skel}(\mu)$ admits a compatible embedding if and only if this conflict graph has maximum degree 2 and either contains no cycle or is a Hamiltonian cycle.

If μ is an R-node, its skeleton is triconnected and therefore has a fixed planar embedding. To test whether the embedding of $\text{skel}(\mu)$ is compatible, we need to check for every bridge B , whether there is a face incident to all its attachments. We consider the graph $\text{skel}'(\mu)$ obtained from $\text{skel}(\mu)$ by subdividing every edge and inserting a vertex into every face that is connected to all incident vertices. We denote the new vertex created by subdividing the virtual edge ε by $v(\varepsilon)$ and the new vertex inserted into the face f by $v(f)$. For a vertex v that already existed in $\text{skel}(\mu)$ we also write $v(v)$. For a bridge B with attachments a_1, \dots, a_k , we need to test whether there is a face f such that for every pair $v(a_i)$ and $v(a_j)$ the path $v(a_i)v(f)v(a_j)$ is contained in $\text{skel}'(\mu)$. To make sure that all paths of length 2 between $v(a_i)$ and $v(a_j)$ include a vertex $v(f)$ corresponding to a face f , we subdivide every edge twice except for those edges incident to a vertex $v(f)$ corresponding to a face.

As $\text{skel}'(\mu)$ is planar, we can use the data structure by Kowalik and Kurowski [16] that can be computed in linear time and supports shortest path queries for constant distance in constant time. More precisely, for any constant d , there exists a data structure that can test in $O(1)$ whether a pair of vertices is connected by a path of length d . If so, a shortest path is returned. We first rule out some easy cases.

If there is an attachment a_i that is a virtual edge in $\text{skel}(\mu)$, the vertex $v(a_i)$ has only four neighbors in $\text{skel}'(\mu)$. Thus, we get the two faces incident to a_i in constant time and can check in $O(k)$ time whether one of them is incident to every attachment of B . We thus assume that all attachments are vertices. If one of these vertices is adjacent to three or more others, there cannot be a compatible embedding. Thus, we either find (in $O(k)$ time) a pair of non-adjacent attachments a_i and a_j or there are only two or three pairwise adjacent attachments. As the latter case is easy, we can assume that we have non-adjacent attachments a_i and a_j . As $\text{skel}(\mu)$ is triconnected, a_i and a_j are incident to at most one common face f . Thus, there is only one path $v(a_i)v(f)v(a_j)$ of length 2 from a_i to a_j in $\text{skel}'(\mu)$, which gives us f in constant time. Then it remains to check whether all other attachments are incident to f , which takes $O(k)$ time. Doing this for every bridge yields a linear-time algorithm testing whether an R-node skeleton admits a compatible embedding. \square

Lemma 19. *Step 2 can be performed in linear time.*

Proof. Recall that Step 2 consist of three parts. First we have to compute the important ①-bridges of S-nodes together with their attachments. For each attachment we then have to test whether it is left- or right-sided. Finally, we have to add the consistency constraints.

As for Step 1, we use Lemma 15, Lemma 16, and Lemma 17 to compute all important ①-bridges together with their attachments. We actually compute these important bridges not only for the S-nodes but also for P- and R-nodes. We use this additional information to compute which attachments are left- and which are right-sided.

Note that the final step of adding the consistency constraints to a global 2-SAT formula is trivial. Thus, it remains to show that we can compute for every attachment whether it is left- or right-sided.

Let μ be a node of the SPQR-tree \mathcal{T} . We iterate over all important bridges of μ (except those we omitted). For every bridge B we iterate over all attachments in μ and for every virtual edge ε among those attachments, we append B to the list $\text{bridges}(\varepsilon)$. Afterwards, $\text{bridges}(\varepsilon)$ contains all important (but not omitted) bridges of μ that have ε as attachment. Note that we can assume that the bridges in $\text{bridges}(\varepsilon)$ are sorted according to an arbitrary but fixed order of the bridges.

Let μ be an S-node with virtual edge ε in $\text{skel}(\mu)$. Let further μ' be the neighboring P- or R-node corresponding to ε and let ε' be the virtual edge in $\text{skel}(\mu')$ corresponding to the S-node μ . For every bridge B in $\text{bridges}(\varepsilon)$ we want to know whether the attachment ε is left- or right-sided. To this end we iterate over the lists $\text{bridges}(\varepsilon)$ and $\text{bridges}(\varepsilon')$ simultaneously. For every bridge B in $\text{bridges}(\varepsilon)$ there are two different cases. Either B also occurs in $\text{bridges}(\varepsilon')$ or it does not. It is not hard to see that the latter can only happen if B is in μ' a pole or parent bridge that was omitted.

If B also occurs in $\text{bridges}(\varepsilon')$, we know from Step 1 that $\text{skel}(\mu')$ has a (unique) face incident to all attachments of B in $\text{skel}(\mu')$. In particular, this face is either the right or the left face of ε' and thus we immediately know whether the attachment ε of B in $\text{skel}(\mu)$ is left- or right-sided.

It remains to consider the case where B occurs in $\text{bridges}(\varepsilon)$ but not in $\text{bridges}(\varepsilon')$. If μ' is the parent of μ , the bridge B must be a pole- or parent bridge in μ' with child attachment ε' . As in the proof of Lemma 18, we can find the (unique) face incident to ε' and one of the poles. As B has to lie in this face we know whether it lies to the right or to the left face of ε' in $\text{skel}(\mu')$ and thus we know whether the attachment ε in $\text{skel}(\mu)$ is left- or right-sided.

If μ' is a child of μ , then B cannot be regular in μ as otherwise B would be semi-regular in μ' and thus contained in $\text{bridges}(\varepsilon')$ (which is not the case we consider). Hence B is either a pole or a parent bridge (recall that semi-regular bridges are also parent bridges). Thus, B is a pole or parent bridge in μ with child attachment ε . By Lemma 13, we can omit all but a constant number of such bridges with ε as child attachment. Thus, we can assume that $\text{bridges}(\varepsilon)$ contains only a constant number of bridges that do not occur in $\text{bridges}(\varepsilon')$. For these bridges we allow a running time linear in the size of $\text{skel}(\mu')$. As μ is the unique parent of μ' this happens only a constant number of times for μ' and thus takes overall linear time.

Let B be such a bridge and let v be the attachment vertex of the common graph G representing the child attachment ε of B in $\text{skel}(\mu)$. We show how to find a child attachment of B in $\text{skel}(\mu')$ in $O(|\text{skel}(\mu')|)$ time. Then we can (as in the cases before) find in constant time which face incident to ε' contains B in $\text{skel}(\mu')$ and we are done. As before we assume to have a DFS-ordering on the leaves of the SPQR-vertex-tree. Then

the leaves that a descendants of an inner node form an interval with respect to this order. These intervals can be easily computed in linear time by processing the SPQR-vertex-tree bottom up once. Afterwards, we can check in constant time whether a vertex v is the descendent of an inner node. Hence we can check in $O(|\text{skel}(\mu')|)$ time which child of μ' is an ancestor of v and thus which virtual edge or vertex in $\text{skel}(\mu')$ represents v . This yields the child attachment of B in $\text{skel}(\mu')$ and we are done. \square

Lemma 20. *Step 3 can be performed in linear time.*

Proof. Let μ be an S-node. Note that every important bridge in μ may alternate with a linear number of important bridges. Thus, there are instances where the planarity constraints have quadratic size. In the following we describe how to compute constraints that are equivalent to the planarity constraints but have linear size. We only consider ①-bridges; for ②-bridges, the same procedure can be applied.

We define the graph H as follows. We start with $H = \text{skel}(\mu)$ and subdivide every edge once. Thus, H has a vertex $v(a)$ for each attachment a in $\text{skel}(\mu)$. For every bridge B with attachments a_1, \dots, a_k , we add a *bridge vertex* $v(B)$ and connect it to the vertices $v(a_1), \dots, v(a_k)$. When using the term *cycle of H* , we refer to the subgraph of H one obtains by removing the bridge vertices.

Assume we have a planar embedding of H . Then, every bridge vertex lies on one of the two sides of the cycle and no two bridges on the same side of the cycle alternate. Conversely, an assignment of the bridges to the two sides of the cycles such that no two bridges alternate yields a planar embedding. The choices the planarity constraints leave are thus equivalent to the embedding choices of H .

As H is biconnected, the embedding choices consist of reordering parallel edges in P-nodes and mirroring R-nodes of the SPQR-tree \mathcal{T}_H of H . Let η be a P-node in \mathcal{T}_H . If the embedding of $\text{skel}(\eta)$ determines on which side of the cycle a bridge B lies, then B is clearly not alternating with any other bridge. For an R-node η of \mathcal{T}_H , fixing the embedding of $\text{skel}(\eta)$ to one of the two flips determines the side for some of the bridges. We create a new binary variable x_η with the interpretation that $x_\eta = 0$ if $\text{skel}(\eta)$ is embedded according to a reference embedding and $x_\eta = 1$ if the embedding is flipped. For a bridge B whose side is determined by the embedding of $\text{skel}(\eta)$ we can then add the constraint $x_\eta = x_B^\mu$ or $x_\eta \neq x_B^\mu$ (depending on whether the reference embedding of $\text{skel}(\eta)$ fixes B to the left or right side of the cycle).

It is not hard to see that one can compute for each R-node η the bridges whose side is determined by the embedding of $\text{skel}(\eta)$ in overall linear time in the size of H . Thus, we get the above constraints (which are equivalent to the planarity constraints) in $O(|H|)$ time. Clearly, the size of H is linear in the size of $\text{skel}(\mu)$ plus the total number of attachments of important bridges in μ . Thus, we get an overall linear running time. \square

Lemma 21. *Step 4 can be performed in linear time.*

Proof. To add the union-bridge constraints, we have to group the ①-bridges that are important in an S-node μ according to their union bridges. To this end, we once create a global array A with one entry $A[B']$ for each union bridge B' (which we can access in

constant time). Consider an ①-bridge B that is important in μ (and was not omitted). Let B' be the union bridge containing B (we can get B' in constant time as every ①-bridge is contained in only one union bridge). If the entry $A[B']$ was not modified while processing μ so far, we clear $A[B']$ (which might contain something from previous nodes) and set $A[B']$ to be a list containing only B . If $A[B']$ was already modified, we append B to the list $A[B']$. We can keep track of which entries of A were already modified by using timestamps.

For every union bridge B' that contains an important ①-bridge, the entry $A[B']$ holds a list of all important ①-bridges of μ that belong to B' . For each of these lists we add the union-bridge constraint for every pair of consecutive ①-bridges. The transitivity enforces all pairwise union-bridge constraints (although not explicitly stated). Clearly, this procedure takes linear time in the number of important ①-bridges. \square

For Step 5, assume the embedding of the biconnected common graph G is fixed. We have to test whether this embedding of G can be extended to a SEFE that satisfies the union-bridge constraints. Thus, we basically have to assign each union bridge to a face of G such that no two ①-bridges and no two ②-bridges alternate.

We first distinguish three different types of union bridges. Let B be a union bridge. A face f of G is **FEASIBLE** for B if all attachment vertices of B are incident to f . We say that B is *flexible* if it has at least three feasible faces. We say that B is *binary* if B has two feasible faces and *fixed* if it has only one feasible face. If there is a bridge that has no feasible face then the instance is obviously not solvable.

The overall strategy for Step 5 is the following. We first determine which union bridges are flexible, binary, and fixed, respectively. We first assign the fixed union bridges to their faces. For the binary union bridges, we can encode the decision for one of the two possible faces with a binary variable. For two union bridges with a common feasible face that include alternating ①-bridges (for the same $i \in \{1, 2\}$), we have to make sure that they are not embedded into the same face. Note that these kind of conditions are very similar to the planarity constraints we had in Step 3, which again leads to a 2-SAT formula. Any solution of this formula induces an assignment of the binary union bridges to faces. Finally, we check whether the flexible union bridges can be added. For this to work we have to show that this final step of assigning the flexible bridges is independent from the solution we chose for the 2-SAT formula. We obtain the following lemma.

Lemma 22. *Step 5 can be performed in linear time.*

Proof. Let B be a flexible union bridge and assume all binary and fixed union bridges are already assigned to faces. We first show the following. If B cannot be embedded into one of its feasible faces (due to alternating ①-bridges), then B cannot be embedded into this face even when omitting all binary union bridges. This shows that the above strategy of first assigning the fixed, then the binary, and finally the flexible union bridges to faces is correct. Afterwards we show how to do it in linear time.

As the union bridge B is flexible, it has at least three feasible faces. As the common graph G is biconnected, B can have only two attachment vertices; let u and v be these attachment vertices. Moreover, u and v must be the poles of a P-node μ of the SPQR-tree

of G . Let $\varepsilon_1, \dots, \varepsilon_k$ be the virtual edges of $\text{skel}(\mu)$ appearing in this order and let f_i be the face between ε_i and ε_{i+1} (subscripts are considered modulo k). The feasible faces of B are exactly the faces f_1, \dots, f_k .

Assume B cannot be assigned to the feasible face f_i , i.e., an ①-bridge contained in B alternates with a ①-bridge belonging to another union bridge B' that was assigned to f_i . As u and v are the only attachment vertices of B , B' must have attachment vertices u' and v' with $u', v' \notin \{u, v\}$ that belong to the expansion graphs of ε_i and ε_{i+1} , respectively. Then u' and v' can share only a single face, namely f_i , and thus B' is a fixed union bridge. This shows the above claim and thus it remains to show how to implement the procedure in linear time.

Let B be an arbitrary union bridge. We show how to detect whether B is flexible, binary, or fixed in $O(|B|)$ time. For B to be flexible, it must have only two different attachment vertices. This can be easily tested in $O(|B|)$ time. If B has only two attachment vertices u and v , we need to test whether u and v are the poles of a P-node in the SPQR-tree of G . We show how this can be done in constant time. To this end, let \mathcal{T} be the SPQR-vertex-tree of G . Assume μ is a P-node with poles u and v . Let η be the parent of μ . Then $\text{skel}(\eta)$ contains both vertices u and v but at most one of them as pole. Assume without loss of generality that u is not a pole of $\text{skel}(\eta)$. Then the leaf u of \mathcal{T} is a child of η and thus we find η in constant time (together with the vertex u in $\text{skel}(\eta)$). If v is not a pole of $\text{skel}(\eta)$, we find v in $\text{skel}(\eta)$ in the same way, otherwise v is a pole and we also get v in $\text{skel}(\eta)$ in constant time (by checking both poles). As $\text{skel}(\eta)$ is a planar graph we can get the virtual edge between the two vertices u and v in constant time via a shortest-path data structure [16]. Thus, we also find the corresponding child μ having u and v as poles. Hence, we can either find the P-node μ with poles u and v in constant time (implying that B is flexible) or we can conclude that such a P-node does not exist (implying that B is not flexible).

Next we determine whether B is binary or fixed. First note that the bridge B (that is not flexible) is binary if and only if there exists an S-node μ such that every attachment vertex of B is a vertex in $\text{skel}(\mu)$. This can be tested in $O(|B|)$ time using the SPQR-vertex-tree. Assume μ is the S-node such that every attachment vertex of B is a vertex in $\text{skel}(\mu)$. We can handle the case where B has only the two poles of $\text{skel}(\mu)$ as attachment vertex analogously to the case above (about flexible bridges) except that μ is an S-node instead of a P-node. Thus, assume that v is an attachment vertex of B that is not a pole of $\text{skel}(\mu)$. Then we can find μ in constant time as it is the parent of the leaf v in \mathcal{T} . Every other attachment vertex in B is either also a child of μ or a pole of $\text{skel}(\mu)$, which we can check in constant time per attachment vertex. Thus, we can detect in $O(|B|)$ time whether B is binary or fixed.

At this point we know which bridges are binary and which are flexible. All remaining bridges are either fixed or do not have a feasible face at all, which implies that there is no SEFE. We show for such a bridge B how we can assign it to its unique feasible face or decide that such a face does not exist. Recall from Lemma 18 that we can compute in constant time a face that is shared by a given pair of vertices (or conclude that such a face does not exist). If B has only two attachment vertices u and v , then we can either find the unique feasible face of B or decide that B has no feasible face in constant time.

We can thus assume that B has at least three attachment vertices. Let u , v , and w be three attachment vertices of B . In constant time, we find a face $f_{u,v}$ that is incident to u and v . Analogously we find faces $f_{u,w}$ and $f_{v,w}$. There are two different cases. If there is a pair of vertices among u , v , and w that shares only a single face, then one of the faces $f_{u,v}$, $f_{u,w}$, or $f_{v,w}$ is the only possible feasible face of B . We can check that in $O(|B|)$ time. Otherwise, assume there is a face $f \notin \{f_{u,v}, f_{u,w}, f_{v,w}\}$ that is feasible for B . Then u and v are commonly incident to at least two different faces (namely f and $f_{u,v}$) and thus $\{u, v\}$ is a separating pair of an edge. The same holds for u and w and for v and w . In this case there must exist a node μ in the SPQR-tree of G such that $\text{skel}(\mu)$ contains the triangle u, v, w . Note that we can find this node as we did before for the flexible and binary bridges.

If μ is an S-node, B must contain another attachment vertex x (otherwise B is binary). Then x is contained in the expansion graph of one of the three virtual edges in $\text{skel}(\mu)$. Assume without loss of generality that x belongs to the expansion graph of uv . Then w and x share only a single face (otherwise w and x would be a separating pair which contradicts the fact that $\text{skel}(\mu)$ is a triangle). Thus, we can find the desired face f by finding a common face of x and w in constant time. Of course one then needs to check if this face is actually incident to each attachment vertex in B (B has no feasibly face if not).

It remains to consider the case that μ is an R-node. First test whether the triangle u, v, w forms a face in $\text{skel}(\mu)$. If so, this face is unique and thus we know the only potentially feasible face of B . Note that this gives us only the face in $\text{skel}(\mu)$. However, one can easily compute a mapping from the faces of skeletons to the faces in the actual graph in linear time in the size of G (this has to be done only once for all bridges).

To conclude, we now ensured that every union bridge that is neither flexible nor binary is fixed and we assigned the fixed union bridges to their unique feasible faces. Let us continue with the binary bridges. For every binary bridge B we already computed the S-node μ containing all the attachment vertices of B . Note that this already gives us the two possible faces in which B can be embedded (of course we again have to translate from faces in a skeleton to faces in G).

When assigning the binary bridges to faces, we have to make sure that no two ①-bridges alternate. This can be ensured using a 2-SAT formula as for Step 3. As before we can compute and solve this 2-SAT formula in linear time. Thus, it remains to add the flexible bridges.

Let μ be a P-node of the SPQR-tree of G and let s and t be the poles of $\text{skel}(\mu)$. Let further $\varepsilon_1, \dots, \varepsilon_k$ be the virtual edges of $\text{skel}(\mu)$ and let the faces f_1, \dots, f_k be defined as before. Assume we still know the important bridges for μ from the previous steps. Assume the union flexible bridge B contains only ①-bridges. Clearly, we can embed B into f_i if and only if there is no ①-bridge with attachments ε_i and ε_{i+1} . The analogous statement holds if B contains only ②-bridges. If B contains both, ①- and ②-bridges, we can embed it into f_i if and only if there is neither a ① nor a ② bridge with attachments ε_i and ε_{i+1} . Thus, we can check in $O(|\text{skel}(\mu)|)$ time for an appropriate face for B . Note that we cannot afford this amount of time for every flexible bridge with attachments s and t . However, consider two such bridges B and B' as equivalent in the sense that

B contains ①- and ②-bridges if and only if B' contains ①- and ②-bridges, respectively. Then B and B' can be assigned to the same face. Thus, we have to spend $O(|\text{skel}(\mu)|)$ time only a constant number of times for each P-node. This concludes the proof. \square

Theorem 7. SEFE with union bridge constraints can be solved in linear time if the common graph is biconnected.

5 Edge Orderings and Relative Positions

In this section, we consider a SEFE instance (G^\oplus, G^\otimes) that has common P-node degree 3 and simultaneous cutvertices of common degree at most 3. Recall that (G^\oplus, G^\otimes) admits a simultaneous embedding if and only if G^\oplus and G^\otimes admit planar embeddings that have consistent edge orderings and consistent relative positions on the common graph. We show how to address both requirements (more or less) separately, by formulating necessary and sufficient constraints using equations and inequalities on Boolean variables. Moreover, we show how to incorporate equations and inequalities equivalent to block-local common-face constraints. Together with the preprocessing algorithms from the previous sections, this leads to a polynomial time algorithm for instances with P-node degree 3 and simultaneous cutvertices of common degree at most 3.

Before we can follow this strategy, we need to address one problem. The relative position of a component H' of G with respect to another connected component H , denoted by $\text{pos}_H(H')$, is the face of H containing H' . However, the set of faces of H depends on the embedding of H . To be able to handle relative positions independently from edge orderings, we need to express the relative positions independently from faces.

This is done in the following section. Afterwards, we show how to enforce consistent edge orderings (Section 5.2), block-local common-face constraints (Section 5.3), and consistent relative positions (Section 5.4). Finally, we conclude in Section 5.5.

Before we start, we need one more definition. Assume we have a set X of binary variables such that every variable $x \in X$ corresponds to a binary embedding choice in a given graph G . Let $\alpha: X \rightarrow \{0, 1\}$ be a *variable assignment*. We say that an embedding of G *realizes* the assignment α , if the embedding decision in G corresponding to a variable $x \in X$ fits to the value $\alpha(x)$. Note that not every variable assignment can be realized as the embedding choices can depend on each other.

5.1 Relative Positions with Respect to a Cycle Basis

A *generalized cycle* C in a graph H is a subset of its edges such that every vertex of H has an even number of incident edges in C . The *sum* $C \oplus C'$ of two generalized cycles is the symmetric difference between the edge sets, i.e., an edge e is contained in $C \oplus C'$ if and only if it is contained in C or in C' but not in both. The resulting edge set $C \oplus C'$ is again a generalized cycle. The set of all generalized cycles in H is a vector space over \mathbb{F}_2 . A basis of this vector space is called *cycle basis* of H .

Instead of considering the relative position $\text{pos}_H(H')$ of a connected component H' with respect to another component H , we choose a cycle basis \mathcal{C} of H and show that the

relative positions of H' with respect to the cycles in \mathcal{C} suffice to uniquely define $\text{pos}_{H'}(H')$, independent from the embedding of H . We assume H to be biconnected. All results can be extended to connected graphs by using a cycle basis for each block.

Let C_0, \dots, C_k be the set of facial cycles with respect to an arbitrary planar embedding of H . The set $\mathcal{C} = \{C_1, \dots, C_k\}$ obtained by removing one of the facial cycles is a cycle basis of G . A cycle basis that can be obtained in this way is called *planar cycle basis*. In the following we assume all cycle bases to be planar cycle bases. Moreover, we consider all cycles to have an arbitrary but fixed orientation. The binary variable $\text{pos}_C(p)$ represents the relative position of a point p with respect to a cycle C , where $\text{pos}_C(p) = 0$ and $\text{pos}_C(p) = 1$ have the interpretation that p lies to the right and left of C , respectively.

Theorem 8. *Let H be a planar graph embedded on the sphere, let p be a point on the sphere, and let $\mathcal{C} = \{C_1, \dots, C_k\}$ be an arbitrary planar cycle basis of H . Then the face containing p is determined by the relative positions $\text{pos}_{C_i}(p)$ for $1 \leq i \leq k$.*

Proof. Let f be a face and let C be the corresponding facial cycle. We assume without loss of generality that $C = C_1 \oplus \dots \oplus C_\ell$ holds. We show that the point p belongs to the face f if and only if $\text{pos}_{C_i}(p) = \text{pos}_{C_i}(f)$ holds for $1 \leq i \leq \ell$. Obviously, if $\text{pos}_{C_i}(p) \neq \text{pos}_{C_i}(f)$ holds for one of the basis cycles C_i , then C_i separates p from f and thus p cannot belong to f .

Conversely, we have to show that there is no point lying on the same sides of the cycles C_i for $1 \leq i \leq \ell$ not belonging to f . To this end we define the *position vector* $\text{pos}(p) = (\text{pos}_{C_1}(p), \dots, \text{pos}_{C_\ell}(p))$ of a point p . We show that there is no point outside f having the same position vector as the points inside f . Consider how the position vector of a point p changes when moving it around. First, all points inside f have the same position vector. Second, when p does not lie in f and crosses an edge e while moving it, then e is either contained in zero or in two of the basis cycles C_1, \dots, C_ℓ . This comes from the facts that $C = C_1 \oplus \dots \oplus C_\ell$ holds, that e is not contained in C and that our cycle basis is planar. In the former case the position vector does not change, in the latter case exactly two values toggle. No matter which case applies, the parity of the number of entries with the value LEFT does not change, that is this number either remains odd or even. Thus, this parity is the same for all points outside of f . Finally, when p moves from inside f to the outside of f (or the other way round), it has to cross an edge e contained in the cycle C . Since e has to be contained in exactly one of the cycles C_1, \dots, C_ℓ , exactly one entry in the position vector $\text{pos}(p)$ changes from LEFT to RIGHT or vice versa. Thus the parity of the number of entries with the value LEFT changes. It follows, that for every point not contained in f this parity differs from the parity of the points in f . Thus also the position vector must differ, which concludes the proof. \square

To represent the relative position of one connected component H with respect to another connected component H' , it thus suffices to consider the relative positions of H with respect to cycles in a cycle basis of H' . However, there is one case for which we have a slightly stronger requirement. To motivate this, consider the following example; see also Figure 16.

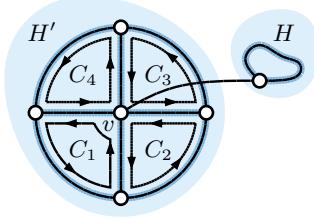


Figure 16. The exclusive edge connecting H to the vertex v of H' requires H to lie to the left of exactly one of the cycles C_1, \dots, C_4 (or to the right of exactly one cycle if the embedding of H' is flipped). This type of constraint cannot be expressed using only equations or inequalities. If, however, the cycle basis contained the facial cycle of the outer face of H' , it would be sufficient to require that H and v lie on the same side of this cycle, which can be expressed using an equation.

Consider the graph G^{\oplus} containing the common graph G with connected components H and H' . Let \mathcal{C} be a cycle basis of H' . Let further v be a vertex of H' that is cutvertex of G^{\oplus} separating H from H' and let $C \in \mathcal{C}$. If v lies to the right of C in a given embedding of G^{\oplus} , then H also lies to the right of C . Conversely, if v lies to the left of C , then H lies to the left of C . However, requiring for every cycle $C \in \mathcal{C}$ (that does not contain v) that v and H lie on the same side of C does not ensure that H lies in a face of H' that is incident to v . Figure 16 shows a somewhat degenerate example, were v is contained in every cycle of \mathcal{C} .

Thus, the relative positions of v with respect to all cycles in \mathcal{C} (that do not contain v) do not uniquely determine a face of $H' - v$. To resolve this issue, we add further cycles of H' to \mathcal{C} . More precisely, an *extended cycle basis* of H' is a set of cycles \mathcal{C} in H' such that \mathcal{C} includes a cycle basis of H' and a cycle basis of $H' - v$ for every vertex v of H' .

Note that one can for example obtain an extended cycle basis of H' as follows. First choose an embedding of H' and start with the corresponding planar cycle basis for \mathcal{C} . For every vertex v , consider the induced embedding of $H' - v$ and add to \mathcal{C} all cycles in the corresponding planar cycle basis of $H' - v$ that are not already contained in \mathcal{C} . It directly follows that an extended cycle basis has $O(n^2)$ size and can be computed in $O(n^2)$ time. Moreover, we get the following lemma.

Lemma 23. *Let H be an embedded planar graph and let \mathcal{C} be an extended cycle basis of H . If a vertex v of H is not incident to f of H , then \mathcal{C} contains a cycle C (not containing v) such that $\text{pos}_C(v) \neq \text{pos}_C(f)$.*

Proof. Consider the graph $H - v$ together with the embedding induced by the embedding of H . Let f_v be the face of $H - v$ that contains v in the embedding of H . As v is not incident to f , we have $f_v \neq f$. By Theorem 8, the cycle basis of $H - v$ (and thus \mathcal{C}) must contain a cycle C such that $\text{pos}_C(f_v) \neq \text{pos}_C(f)$. \square

If we refer to a cycle basis in one of the following sections, we always assume to actually have an extended cycle basis.

5.2 Consistent Edge Orderings

We first assume that the graphs G^{\circledcirc} and G^{\circledast} are biconnected and then show how to extend our approach to exclusive cutvertices and simultaneous cutvertices of common degree 3.

5.2.1 Biconnected Graphs

Let G^{\circledcirc} and G^{\circledast} be biconnected planar graphs. There exists an instance of SIMULTANEOUS PQ-ORDERING that has a solution if and only if G^{\circledcirc} and G^{\circledast} admit embeddings with consistent edge ordering [4]. This solution is based on the *PQ-embedding representation*, an instance of SIMULTANEOUS PQ-ORDERING representing all embeddings of a biconnected planar graph. We describe this embedding representation and show how to simplify it for instances that have common P-node degree 3.

For each vertex v^{\circledcirc} of G^{\circledcirc} , the PQ-embedding representation, denoted by $D(G^{\circledcirc})$, contains the *embedding tree* $T(v^{\circledcirc})$ having a leaf for each edge incident to v^{\circledcirc} , representing all possible orders of edges around v^{\circledcirc} . For every P-node μ^{\circledcirc} in the SPQR-tree $\mathcal{T}(G^{\circledcirc})$ that contains v^{\circledcirc} in $\text{skel}(\mu^{\circledcirc})$ there is a P-node in $T(v^{\circledcirc})$ representing the choice to order the virtual edges in $\text{skel}(\mu^{\circledcirc})$. Similarly, for every R-node μ^{\circledcirc} of G^{\circledcirc} containing v^{\circledcirc} in its skeleton, there is a Q-node in $T(v^{\circledcirc})$ whose flip corresponds to the flip of $\text{skel}(\mu^{\circledcirc})$. As the orders of edges around different vertices of G^{\circledcirc} cannot be chosen independently from each other, so called *consistency trees* are added as common children to enforce Q-nodes stemming from the same R-node in $\mathcal{T}(G^{\circledcirc})$ to have the same flip and P-nodes stemming from the same P-node to have consistent (i.e., opposite) orders. Every solution of the resulting instance corresponds to a planar embedding of G^{\circledcirc} and vice versa [4].

As we are only interested in the order of common edges, we modify $D(G^{\circledcirc})$ by projecting each PQ-tree to the leaves representing common edges. As G^{\circledcirc} and G^{\circledast} have common P-node degree 3, all P-nodes of the resulting PQ-trees have degree 3 and can be assumed to be Q-nodes representing a binary decision. We call the resulting instance *Q-embedding representation* and denote it by $D(G^{\circledcirc})$.

Let μ^{\circledcirc} be an R-node of the SPQR-tree $\mathcal{T}(G^{\circledcirc})$ whose embedding influences the ordering of common edges around a vertex. Then the Q-embedding representation contains a consistency tree consisting of a single Q-node representing the flip of $\text{skel}(\mu^{\circledcirc})$. We associate the binary variable $\text{ord}(\mu^{\circledcirc})$ with this decision.

For a P-node μ^{\circledcirc} we get a similar result. Let u^{\circledcirc} and v^{\circledcirc} be the poles of μ^{\circledcirc} . If the consistency tree enforcing a consistent decision in the embedding trees $T(u^{\circledcirc})$ and $T(v^{\circledcirc})$ has degree 3, its flip represents the embedding decision for $\text{skel}(\mu^{\circledcirc})$ and we again get a binary variable $\text{ord}(\mu^{\circledcirc})$. Otherwise, this consistency tree contains two or less leaves and can be ignored. Then the choices for the Q-nodes corresponding to μ^{\circledcirc} in $T(u^{\circledcirc})$ and $T(v^{\circledcirc})$ are independent and we get one binary variable for each of these Q-nodes. We denote these variables by $\text{ord}(\mu_u^{\circledcirc})$ and $\text{ord}(\mu_v^{\circledcirc})$.

We call these variables we get for G^{\circledcirc} the \circledcirc -PR-ordering variables. The \circledast -PR-ordering variables are defined analogously. The PR-ordering variables are the union of these two variable sets. Let X^{\circledcirc} be the \circledcirc -PR-ordering variables and assume we have a fixed variable assignment α^{\circledcirc} . Then this variable assignment already determines all edge orderings of

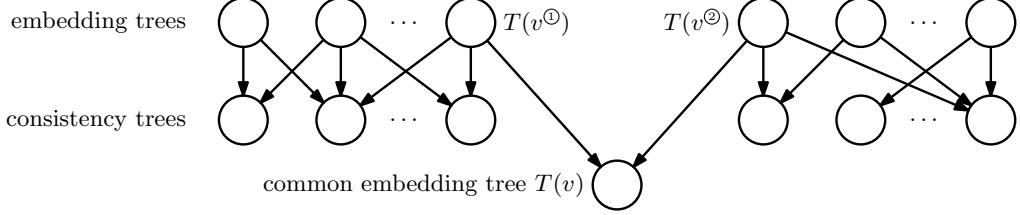


Figure 17. The Q-embedding representations of G^1 and G^2 together with the common embedding tree $T(v)$ of v .

the common graph, i.e., every embedding of G^1 realizing the assignment α^1 induces the same edge orderings on G . In the following we describe a set of necessary equations and inequalities on the PR-ordering variables that ensure that G^1 and G^2 induce the same edge orderings on G .

For a common vertex v occurring as v^1 and v^2 in G^1 and G^2 , respectively, we add a so-called *common embedding tree* $T(v)$ (consisting of a single P-node) as child of the embedding trees $T(v^1)$ and $T(v^2)$ in the Q-embedding representations of G^1 and G^2 ; see Figure 17. Obviously, this common child ensures that the common edges around v are ordered the same with respect to G^1 and G^2 . Adding the common embedding tree for every common vertex yields the instance $D(G^1, G^2)$.

As the embedding trees (which contain only Q-nodes) are the sources of $D(G^1, G^2)$, normalizing $D(G^1, G^2)$ yields an equivalent instances containing no P-nodes [4]. Instances with this property are equivalent to a set of Boolean equations and inequalities, containing one variable for each Q-node in each PQ-tree (and thus includes the PR-ordering variables). We call this set of equations and inequalities the *biconnected PR-ordering constraints*. To obtain the following lemma, it remains to proof the running time.

Lemma 24. *Let G^1 and G^2 be two biconnected graphs with common P-node degree 3 and let α be a variable assignment for the PR-ordering variables. The graphs G^1 and G^2 admit embeddings that realize α and have consistent edge orderings if and only if α satisfies the biconnected PR-ordering constraints.*

The biconnected PR-ordering constraints have size $O(n)$ and can be computed in $O(n)$ time.

Proof. Let $D(G^1, G^2)$ be the instances of SIMULTANEOUS PQ-ORDERING as described above. Clearly, $D(G^1, G^2)$ can be constructed in linear time and its size is linear in the size of the input graphs. In general, the equations and inequalities for a given instance of SIMULTANEOUS PQ-ORDERING can be computed in quadratic time [4]. In this specific case, it can be done in linear time for the following reasons.

For every arc in $D(G^1, G^2)$ one needs to compute the normalization of the child (which takes linear time in the size of the parent [4]) and a mapping from each inner node of the parent to its representative in the child (which takes again linear time in the size of the parent [4]). When computing the Q-embedding representations for G^1 and G^2 from their SPQR-trees (which can be done in linear time), we can make sure that the resulting

instances are already normalized and that we already know the mapping from the nodes of the embedding trees to the consistency trees. The remaining arcs in $D(G^\oplus, G^\otimes)$ are arcs from embedding trees to common embedding trees. Thus, for every PQ-tree T in $D(G^\oplus, G^\otimes)$, it suffices to normalize a single outgoing edge, which can be done in time linear in the size of T . \square

5.2.2 Allowing Cutvertices

In the following, we extend this result to the case where we allow exclusive cutvertices and simultaneous cutvertices of common degree 3. Let $B_1^\oplus, \dots, B_k^\oplus$ be the blocks of G^\oplus and let $B_1^\otimes, \dots, B_\ell^\otimes$ be the blocks of G^\otimes . We say that embeddings of these blocks have *blockwise consistent edge orderings* if for every pair of blocks B_i^\oplus and B_j^\otimes sharing a vertex v the edges incident to v they share are ordered consistently. To have consistent edge orderings, it is obviously necessary to have blockwise consistent edge orderings.

When composing the embeddings of two blocks that share a cutvertex, the edges of each of the two blocks have to appear consecutively (note that this is no longer true for three or more blocks), which leads to another necessary condition. Let v be an exclusive cutvertex of G^\oplus . Then v is contained in a single block of G^\otimes whose embedding induces an order O^\otimes on all common edges incident to v . Let B_i^\oplus and B_j^\oplus (for $i, j \in \{1, \dots, k\}$ with $i \neq j$) be a pair of blocks containing v and let $O_{i,j}^\otimes$ be the order obtained by restricting O^\otimes to the common edges in B_i^\oplus and B_j^\oplus . Then the common edges of B_i^\oplus must be consecutive in the order $O_{i,j}^\otimes$. If this is true for every pair of blocks at every exclusive cutvertex, we say that the embeddings have *pairwise consecutive blocks*.

Lemma 25. *Two graphs without simultaneous cutvertices admit embeddings with consistent edge orderings if and only if their blocks admit embeddings that have blockwise consistent edge orderings and pairwise consecutive blocks.*

Proof. Let v be a cutvertex in G^\oplus and let $B_1^\oplus, \dots, B_k^\oplus$ be the blocks containing v . Moreover, let O^\otimes be the order of common edges around v given by the unique block of G^\otimes containing v . We only have to show that the embeddings of the blocks $B_1^\oplus, \dots, B_k^\oplus$ can be composed such that they induce the order O^\otimes on the common edges. For $k = 2$ this is clear, as we can choose arbitrary outer faces for B_1^\oplus and B_2^\oplus and combine their embeddings by gluing them together at v . For $k > 2$ it is easy to see that there must be one block, without loss of generality B_1^\oplus , whose edges appear consecutively in O^\otimes . The embeddings of all remaining blocks $B_2^\oplus, \dots, B_k^\oplus$ can be composed by induction such that the edges they contain are ordered the same as in O^\otimes . Moreover, composing the embedding of B_1^\oplus with the resulting embedding of $B_2^\oplus, \dots, B_k^\oplus$ works the same as the composition of two blocks. \square

To extend Lemma 24 to the case where we allow exclusive cutvertices, we consider the Q-embedding representations of each block. Ensuring blockwise consistent edge orderings works more or less the same as ensuring consistent edge orderings in the biconnected case. Moreover, we will see how to add additional PQ-trees to ensure pairwise consecutive blocks.

Note that the Q-embedding representations again yield PR-ordering variables. Fixing these variables determines the edge orderings of common edges in each block of G^{\circledcirc} and in each block of G^{\circledast} . If we have no simultaneous cutvertices, every vertex is either not a cutvertex in G^{\circledcirc} or not a cutvertex in G^{\circledast} . Thus, the PR-ordering variables actually determine all edge orderings of the common graph. Thus, although there are new types of embedding choices at the cutvertices in G^{\circledcirc} and at the cutvertices in G^{\circledast} , these choices are already covered by the PR-ordering variables (at least in terms of edge orderings).

Let us formally describe the instance of SIMULTANEOUS PQ-ORDERING announced above. Let $B_1^{\circledcirc}, \dots, B_k^{\circledcirc}$ be the blocks of G^{\circledcirc} and let $B_1^{\circledast}, \dots, B_\ell^{\circledast}$ be the blocks of G^{\circledast} . We start with an instance of SIMULTANEOUS PQ-ORDERING containing the Q-embedding representation of each of these blocks. Let v be a vertex of G that is not a cutvertex. Then v is contained in a single block of G^{\circledcirc} and in a single block of G^{\circledast} , let v^{\circledcirc} and v^{\circledast} be the occurrences of v in these blocks. As before, there are two embedding trees $T(v^{\circledcirc})$ and $T(v^{\circledast})$ describing the order of edges around v in G^{\circledcirc} and G^{\circledast} , respectively. As before we can enforce consistent ordering around v by inserting a common embedding tree as a common child of $T(v^{\circledcirc})$ and $T(v^{\circledast})$.

Let v be a cutvertex of G^{\circledcirc} (the case that v is a cutvertex of G^{\circledast} is symmetric). Then v occurs in several blocks of G^{\circledcirc} , without loss of generality $B_1^{\circledcirc}, \dots, B_r^{\circledcirc}$. We denote the occurrences of v in these blocks by $v_1^{\circledcirc}, \dots, v_r^{\circledcirc}$. As v is not a simultaneous cutvertex, it occurs in a single block B^{\circledast} of G^{\circledast} . We denote this occurrence by v^{\circledast} . The embedding tree $T(v^{\circledast})$ contains a leaf for each of common edge incident to v , thus fixing the order of $T(v^{\circledast})$ already fixes the order of all common edges around v . To ensure consistency, we have to enforce that conditions of Lemma 25, i.e., blockwise consistent edge orderings and pairwise consecutive blocks.

Ensuring blockwise consistent edge orderings is equivalent to enforcing that the common edges incident to v in B_i^{\circledcirc} (for $i = 1, \dots, k$) are ordered the same with respect to the Q-embedding representations of B_i^{\circledcirc} and B^{\circledast} . This can be done by inserting a new child consisting of a single P-node as common child of $T(v^{\circledast})$ and $T(v_i^{\circledcirc})$. We call this tree the *blockwise common embedding tree*.

To ensure pairwise consecutive blocks, consider the two blocks B_i^{\circledcirc} and B_j^{\circledcirc} . We create a PQ-tree $T_{i,j}(v^{\circledcirc})$ that has a leaf for each common edge incident to v^{\circledcirc} that belongs to one of the two blocks B_i^{\circledcirc} or B_j^{\circledcirc} . The structure of $T_{i,j}(v^{\circledcirc})$ is chosen such that all common edges in B_i^{\circledcirc} are consecutive; see Figure 18. We call $T_{i,j}(v^{\circledcirc})$ *pairwise consecutivity trees* and add it as a child of the embedding tree $T(v^{\circledast})$, which has a leaf for every common edge incident to v .

We denote the resulting instance of SIMULTANEOUS PQ-ORDERING by $D(G^{\circledcirc}, G^{\circledast})$. As before, all sources in $D(G^{\circledcirc}, G^{\circledast})$ contain only Q-nodes. Thus, normalizing $D(G^{\circledcirc}, G^{\circledast})$ leads to an instance containing only Q-nodes [4], which is again equivalent to a set of equations and inequalities. We call this set the *PR-ordering constraints*.

So far, the PR-ordering constraints only ensure blockwise consistent edge orderings and pairwise consecutive blocks if every cutvertex is an exclusive cutvertex. Recall that there is no need for handling union cutvertices; see Theorem 1. Assume we allow simultaneous cutvertices of common degree 3 and let v be such a cutvertex. Then there are two

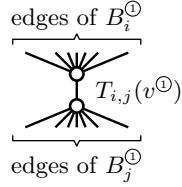


Figure 18. A pairwise consecutivity tree $T_{i,j}(v^{\circledcirc})$.

possibilities. If v does not separate the three common edges incident to v in one of the graphs G^{\circledcirc} (for $i \in \{1, 2\}$), then the PR-ordering variables of G^{\circledcirc} also determine the common edge ordering around v and thus this simultaneous cutvertex actually behaves like an exclusive cutvertex. Otherwise, the common edges incident to v are separated by v in G^{\circledcirc} and in G^{\circledcirc} . Thus, changing the edge ordering of the common edges at v in an embedding of G^{\circledcirc} has no effect on any other edge ordering. As the same holds for G^{\circledcirc} , we actually choose an arbitrary edge ordering for the common edges incident to v , independent from all other edge orderings.

Thus, we do not need to add additional constraints for the case that we allow simultaneous cutvertices of common degree 3. To obtain the following lemma, it remains to prove the running time.

Lemma 26. *Let G^{\circledcirc} and G^{\circledcirc} have common P-node degree 3 and simultaneous cutvertices of common degree at most 3. Let α be a variable assignment for the PR-ordering variables. The graphs G^{\circledcirc} and G^{\circledcirc} admit embeddings that realize α and have consistent edge orderings if and only if α satisfies the PR-ordering constraints.*

The PR-ordering constraints have size $O(n^2)$ and can be computed in $O(n^2)$ time.

Proof. The size of the PR-ordering constraints is linear in the size of the instance $D(G^{\circledcirc}, G^{\circledcirc})$. Clearly, the Q-embedding representations of the blocks have overall linear size. The common embedding tree for a vertex v has size $O(\deg(v))$. Similarly, the blockwise common embedding trees of a vertex v have total size $O(\deg(v))$. However, if G^{\circledcirc} has a linear number of blocks incident to a vertex v , then we get a quadratic number of pairwise consecutivity trees.

To get the PR-ordering constraints from the instance $D(G^{\circledcirc}, G^{\circledcirc})$ we have to compute for each arc in $D(G^{\circledcirc}, G^{\circledcirc})$ the normalization and the mapping from the Q-nodes of the source to their representatives in the target. As in the proof of Lemma 24, the arcs of $D(G^{\circledcirc}, G^{\circledcirc})$ belonging to the Q-embedding representation can be computed in linear time. All remaining arcs have an embedding tree as source. An embedding tree of a vertex v has only $O(\deg(v))$ common embedding trees or blockwise common embedding trees as children. Processing all arcs to these children takes $O(\deg(v)^2)$ time and thus overall $O(n^2)$ time.

It remains to deal with arcs of the following type. The source is the embedding tree $T(v^{\circledcirc})$ and the target is the pairwise consecutivity tree $T_{i,j}(v^{\circledcirc})$ for a pair of blocks $B_i^{\circledcirc} \neq B_j^{\circledcirc}$. Normalizing such an arc would usually take $O(\deg(v))$ time, which has to be done for $O(\deg(v)^2)$ pairwise consecutivity trees, resulting in the running time

$O(\deg(v)^3)$. To improve this, we can make use of the fact that the subtree of $T_{i,j}(v^\circledcirc)$ (after the normalization) that represents only the edges in B_i^\circledcirc has always the same structure, independent of the other block B_j^\circledcirc .

We first compute for each block B_i^\circledcirc the reduction of $T(v^\circledcirc)$ with the leaves belonging to B_i^\circledcirc , which takes $O(\deg(v))$ time for each of the $O(\deg(v))$ blocks. The resulting tree contains a single node η_i separating the leaves belonging to B_i^\circledcirc from all other leaves. We project this tree to the leaves belonging to B_i^\circledcirc to obtain the tree $T_i(v^\circledcirc)$ with root η_i . Computing these trees $T_i(v^\circledcirc)$ together with the mapping from the nodes in $T(v^\circledcirc)$ to their representatives in $T_i(v^\circledcirc)$ takes $O(\deg(v))$ time for each block and thus overall $O(\deg(v)^2)$ time.

The desired (normalized) pairwise consecutivity tree $T_{i,j}(v^\circledcirc)$ can be obtained by identifying the roots η_i and η_j of the trees $T_i(v^\circledcirc)$ and $T_j(v^\circledcirc)$ with each other. Extending the mapping to the resulting tree $T_{i,j}(v^\circledcirc)$ can be easily done in linear time in the size of $T_{i,j}(v^\circledcirc)$. Hence, for the whole instance, we get the running time $O(n^2)$. \square

5.2.3 Cutvertex-Ordering Variables

Let v be an exclusive cutvertex of G^\circledcirc and let $B_1^\circledcirc, \dots, B_k^\circledcirc$ be the blocks incident to v that include common edges incident to v . As mentioned before, the choice of how these blocks are ordered around v and how they are nested into each other is determined by the PR-ordering variables of G^\circledcirc . Consider a pair of blocks B_i^\circledcirc and B_j^\circledcirc . As the choice of how B_i^\circledcirc and B_j^\circledcirc are embedded into each other at v may also have an effect on some relative positions, we would like to have more direct access to this information (not only via the PR-ordering variables of G^\circledcirc).

Let $B^\circledcirc \in \{B_1^\circledcirc, \dots, B_k^\circledcirc\}$ be a block of G^\circledcirc that contains the common edge e incident to v and let e_1 and e_2 be two common edges incident to v that are contained in one block distinct from B^\circledcirc . We create a variable $\text{ord}(e_1, e_2, B^\circledcirc)$ to represent the binary decision of ordering the edges e_1 , e_2 , and e in this order or in its reversed order. Note that this is independent from the choice of the edge e of B^\circledcirc (the blocks are pairwise consecutive in every embedding). We create such a variable for every such triple e_1 , e_2 , and B^\circledcirc and call them the *exclusive cutvertex-ordering variables*.

To make sure that the exclusive cutvertex-ordering variables are consistent with the PR-ordering variables, it suffices to slightly change the above instance $D(G^\circledcirc, G^\circledcirc)$ of SIMULTANEOUS PQ-ORDERING. Let $\text{ord}(e_1, e_2, B^\circledcirc)$ be an exclusive cutvertex-ordering variable for the cutvertex v . Let e be a common edge in B^\circledcirc incident to v and let v^\circledcirc be the occurrence of v in G^\circledcirc . Then $D(G^\circledcirc, G^\circledcirc)$ contains the embedding tree $T(v^\circledcirc)$ that has a leaf for every common edge incident to v . This includes e_1 , e_2 , and e . We create a new PQ-tree $T(e_1, e_2, e)$ with three leaves corresponding to e_1 , e_2 , and e and add it as child of $T(v^\circledcirc)$. In every solution of the resulting instance of SIMULTANEOUS PQ-ORDERING, the value of $\text{ord}(e_1, e_2, B^\circledcirc)$ then simply corresponds to the orientation chosen for PQ-tree $T(e_1, e_2, e)$.

Adding this tree for every exclusive cutvertex-ordering variable establishes the desired connection between these variables and the PR-ordering variables. We call the constraints

we get from the resulting instance of SIMULTANEOUS PQ-ORDERING in addition to the PR-ordering constraints the *cutvertex-ordering constraints*.

Let v be a simultaneous cutvertex of common degree 3 such that the common edges incident to v are separated by v in G^{\circledcirc} and in G^{\circledast} . Recall that this is the unique case, where PR-ordering variables do not determine the order of the common edges around v . Let e_1 , e_2 , and e_3 be the common edges incident to v . To make sure that assigning values to all variables actually determines all edge-orderings, we add the variable $\text{ord}(e_1, e_2, e_3)$ associated with the order of these three edges. Recall that changing this order in G^{\circledcirc} or G^{\circledast} has no effect on any other edge ordering. Hence, there is no need to add further constraints. If two of the edges, without loss of generality e_1 and e_2 , belong to the same block of G^{\circledcirc} and e_3 belongs to another block B^{\circledcirc} , we denote $\text{ord}(e_1, e_2, e_3)$ also by $\text{ord}(e_1, e_2, B^{\circledcirc})$ to obtain consistency with the naming of the exclusive cutvertex-ordering variables.

We call these variables together with the exclusive cutvertex-ordering variables the *cutvertex-ordering variables*. The PR-ordering variables together with the cutvertex-ordering variables are simply called *ordering variables*. Moreover, the PR-ordering constraints together with the cutvertex-ordering constraints are called *ordering constraints*. To extend Lemma 26 to incorporate the cutvertex-ordering variables, it remains to show that the cutvertex-ordering constraints have $O(n^3)$ size and can be computed in $O(n^3)$ time.

Lemma 27. *Let G^{\circledcirc} and G^{\circledast} have common P-node degree 3 and simultaneous cutvertices of common degree at most 3. Let α be a variable assignment for the ordering variables. The graphs G^{\circledcirc} and G^{\circledast} admit embeddings that realize α and have consistent edge orderings if and only if α satisfies the ordering constraints.*

The ordering constraints have size $O(n^3)$ and can be computed in $O(n^3)$ time.

Proof. The size of the cutvertex-ordering constraints is clearly linear in the number of cutvertex-ordering variables. For each cutvertex v , the number of cutvertex-ordering variables is clearly in $O(\deg(v)^3)$. Thus, it remains to show how to get the cutvertex-ordering constraints from the resulting instance $D(G^{\circledcirc}, G^{\circledast})$ of SIMULTANEOUS PQ-ORDERING.

Let v be a cutvertex in G^{\circledcirc} and let v^{\circledast} be the occurrence of v in G^{\circledast} . For every cutvertex-ordering variable of v we have three common edges e_1 , e_2 , and e and we need to find the node η of the embedding tree $T(v^{\circledast})$ that separates the leaves corresponding to e_1 , e_2 , and e from each other. When rooting $T(v^{\circledast})$ at e , this node η is the lowest common ancestor of e_1 and e_2 . Thus, after $O(\deg(v))$ preprocessing time, we can get the cutvertex-ordering constraint for every cutvertex-ordering variable that includes e in constant time per variable [14]. We have to spend this $O(\deg(v))$ preprocessing time at most once for each common edge incident to v , yielding a total preprocessing time of $O(\deg(v)^2)$. As we have $O(\deg(v)^3)$ cutvertex-ordering variables for v , the running time is dominated by the constant time LCA-queries, which yield the running time $O(\deg(v)^3)$. For the whole instance, this gives the claimed $O(n^3)$ running time. \square

5.3 Common-Face Constraints

Recall from Section 3.4 that we can assume that there are no bridges that are block-local and exclusive one-attached if we in return solve SEFE with block-local common-face constraints. In this section, we show how to handle these additional constraints. To this end, we show that satisfying block-local common-face constraints in a given instance of SEFE is equivalent to satisfying a set of equations and inequalities. The union of these constraints with the constraints from the previous section thus enforce that the embeddings of each common connected component are consistent and satisfy given block-local common-face constraints.

Let B be a block of the common graph. Let μ be an R-node of B . Then we introduce the binary variable $\text{ord}(\mu)$ where $\text{ord}(\mu) = 0$ indicates that $\text{skel}(\mu)$ is embedded according to its reference embedding and $\text{ord}(\mu) = 1$ indicates that $\text{skel}(\mu)$ is flipped. In case μ is a P-node of B , we can assume by Theorem 2 that the union-link graph of μ is connected. Recall that this implies that the embedding of $\text{skel}(\mu)$ is fixed up to a flip (Lemma 2). Thus, we also get a reference embedding for $\text{skel}(\mu)$ and can describe the embedding choice for $\text{skel}(\mu)$ with a binary variable $\text{ord}(\mu)$. We call these variables the *common PR-node variables*.

It follows directly from Section 4.3 that common-face constraints for B are equivalent to a set of equations and inequalities on the variables $\text{ord}(\mu)$ (we basically get the consistency and union-bridge constraints from Step 2 and Step 4).

Note that the constraints from Section 5.2 enforcing consistent edge orderings do not contain common PR-node variables. They only contain PR-ordering variables determining the embeddings of G^{\circledR} and G^{\circledL} (Lemma 26). As fixing the PR-node variables for G^{\circledR} fixes the embedding of G , it also fixes the values for all common PR-node variables. It remains to show that this dependency of the common PR-node variables from the PR-ordering variables in G^{\circledR} can be expressed using a set of equations and inequalities.

To this end, consider a common vertex v in the common block B and let B^{\circledR} be the block of G^{\circledR} containing B . Let $T(v)$ be the embedding tree of v in B , i.e., the PQ-tree describing the possible edge orderings of common edges incident to v . Note that each inner node of $T(v)$ is actually a Q-node and that $T(v)$ has one inner node for each P- or R-node whose embedding affects the edge ordering around v . Let v^{\circledR} be the occurrence of v in B^{\circledR} and let $T(v^{\circledR})$ be the embedding tree of v^{\circledR} in B^{\circledR} projected to the common edges incident to v . Then $T(v^{\circledR})$ describes all orders of common edges incident to v that can be induced by an embedding of B^{\circledR} .

Clearly, $T(v^{\circledR})$ is more restrictive than $T(v)$ in the sense that every order represented by $T(v^{\circledR})$ is also represented by $T(v)$. Thus, $T(v^{\circledR})$ is a reduction of $T(v)$. It is not hard to see [4] that every Q-node in $T(v)$ has a unique Q-node in $T(v^{\circledR})$ (called its *representative*) that determines its flip. Thus, for every P- or R-node μ of G , we find at least one vertex v such that $\text{ord}(\mu)$ corresponds to the flip of a Q-node in $T(v)$, which corresponds to a flip of a Q-node in $T(v^{\circledR})$, which corresponds to a PR-ordering variable $\text{ord}(\mu^{\circledR})$ (or $\text{ord}(\mu_v^{\circledR})$) of G^{\circledR} . Thus, $\text{ord}(\mu) = \text{ord}(\mu^{\circledR})$ or $\text{ord}(\mu) \neq \text{ord}(\mu^{\circledR})$ gives us the desired connection between the common PR-node variables and the PR-ordering variables.

We call the set of all equations and inequalities described in this section the *common-face*

constraints. With the results from Section 4.3 (and with standard PQ-tree operations), we can compute the common-face constraints in linear time. This yields the following lemma where n is the total input size, i.e., the size of the two graphs plus the size of the common-face constraints.

Lemma 28. *Let $(G^{\circledR}, G^{\circledast})$ be an instance of SEFE with common-face constraints and let α be a variable assignment for the PR-ordering and the common PR-node variables. Every embedding of G^{\circledR} realizing α satisfies the common-face constraints if and only if α satisfies the common-face constraints.*

The common-face constraints have $O(n)$ size and can be computed in $O(n)$ time.

5.4 Consistent Relative Positions

Let H and H' be two connected components of the common graph G . To represent the relative position $\text{pos}_{H'}(H)$ of H with respect to H' , we use the relative positions $\text{pos}_C(H)$ of H with respect to the cycles C in an extended cycle basis of H' . With $\text{pos}_C(H) = 0$ and $\text{pos}_C(H) = 1$, we associate the cases that H lies to the right of C and to the left of C , respectively. We call these variables the *component position variables*. Note that fixing all position variables determines all relative positions of common connected components with respect to each other (Theorem 8). Thus, fixing the PR-ordering variables (which also fixes the cutvertex-ordering variables) and the position variables completely determines the embedding of the common graph G .

In this section, we give a set of necessary equations and inequalities on the position variables of two graphs G^{\circledR} and G^{\circledast} that enforce consistent relative positions on their common graph G . As fixing the PR-ordering variables may also determine some position variables, we also have to make sure that these two types of variables are consistent with each other.

Let G^{\circledR} be a connected planar graph containing G , let C be a cycle in G (a cycle from the extended cycle basis), and let H be a connected component of G not containing C . Depending on how C and H are located in G^{\circledR} , different embedding choices of G^{\circledR} determine the relative position $\text{pos}_C(H)$ [7]. We quickly list these embedding choices here and describe the constraints arising from them in the following sections.

Let μ^{\circledR} be an R-node of G^{\circledR} such that C induces a cycle κ in $\text{skel}(\mu^{\circledR})$. If $\text{skel}(\mu^{\circledR})$ has a virtual edge ε that is not part of κ such that the expansion graph $\exp(\varepsilon)$ includes a vertex of H , then the embedding of $\text{skel}(\mu^{\circledR})$ determines the relative position $\text{pos}_C(H)$. We say that $\text{pos}_C(H)$ is *determined by the R-node μ^{\circledR}* .

Let μ^{\circledR} be a P-node of G^{\circledR} such that C induces a cycle κ in $\text{skel}(\mu^{\circledR})$. Then κ consists of two virtual edges ε_1 and ε_2 . The relative position $\text{pos}_C(H)$ is determined by the embedding of $\text{skel}(\mu^{\circledR})$ if H is contained in the expansion graph of a virtual edge ε with $\varepsilon \neq \varepsilon_i$ for $i \in \{1, 2\}$. We say that $\text{pos}_C(H)$ is *determined by the P-node μ^{\circledR}* .

If C and H belong to the same block B^{\circledR} of G^{\circledR} , then $\text{pos}_C(H)$ is either determined by an R-node or by a P-node of B^{\circledR} [7]. Otherwise, let v be the cutvertex in G^{\circledR} that separates C from H and belongs to the block of C . If v is not contained in C , then we introduce the variable $\text{pos}_C(v)$ corresponding to the decision of embedding v to the right

or to the left of C . We call such a variable the *cutvertex position variables*. Clearly, H and v lie on the same side in each embedding of G^\oplus . Moreover, the relative position of v with respect to C is determined by an R-node or by a P-node μ^\oplus (the belong to the same block of G^\oplus). In this case, we also say that both variables, $\text{pos}_C(v)$ and $\text{pos}_C(H)$, are determined by the R-node or P-node μ^\oplus . Moreover, we also say that $\text{pos}_C(H)$ is *determined by* $\text{pos}_C(v)$.

If the cutvertex v is contained in C , then the relative position $\text{pos}_C(H)$ is determined by the embedding choices made at the cutvertex. We distinguish two cases. Let S^\oplus be the split component with respect to the cutvertex v that contains H . If S^\oplus includes a common edge incident to v , we say that $\text{pos}_C(H)$ is *determined by the common cutvertex v* . Otherwise, we say that $\text{pos}_C(H)$ is *determined by the exclusive cutvertex v* (note that v might still be a cutvertex of the common graph in this case). These two cases are in so far different as changing $\text{pos}_C(H)$ affects the edge ordering of the common graph in the former case, whereas it does not in the latter case.

The component position variables together with the cutvertex position variables are simply called *position variables*. In the following sections, we describe for each of the four cases different constraints in the form of equations and inequalities on position variables, PR-ordering variables, and cutvertex-ordering variables.

5.4.1 Relative Positions Determined by R-Nodes

We start with the simplest case, where $\text{pos}_C(H)$ is determined by an R-node μ^\oplus of G^\oplus . If $\text{pos}_C(H)$ is also determined by a cutvertex position variable $\text{pos}_C(v)$, we simply set $\text{pos}_C(H) = \text{pos}_C(v)$ to make sure that the cutvertex v and the component H are on the same side of C .

Otherwise, C induces a cycle κ in $\text{skel}(\mu^\oplus)$ and H shares a vertex with the expansion graph of a virtual edge ε not belonging to κ . The PR-ordering variable $\text{ord}(\mu^\oplus)$ determines whether $\text{skel}(\mu^\oplus)$ is embedded according to its reference embedding ($\text{ord}(\mu^\oplus) = 0$) or whether it is flipped ($\text{ord}(\mu^\oplus) = 1$).

Assume ε lies to the right of κ in the reference embedding of $\text{skel}(\mu^\oplus)$. Then $\text{ord}(\mu^\oplus) = 0$ implies that ε lies to the right of κ , which implies that H lies to the right of C , i.e., $\text{pos}_C(H) = 0$. Moreover, flipping $\text{skel}(\mu^\oplus)$ brings ε to the left of κ . Thus, $\text{ord}(\mu^\oplus) = 1$ implies $\text{pos}_C(H) = 1$, which yields $\text{ord}(\mu^\oplus) = \text{pos}_C(H)$ as necessary condition. If ε lies to the left of κ in the reference embedding, we obtain $\text{ord}(\mu^\oplus) \neq \text{pos}_C(H)$ instead.

Analogously, we set $\text{ord}(\mu^\oplus) = \text{pos}_C(v)$ or $\text{ord}(\mu^\oplus) \neq \text{pos}_C(v)$ for every cutvertex position variable $\text{pos}_C(v)$ determined by μ^\oplus .

Sufficiency. We call the constraints defined in this section the *R-node constraints*. The following lemma states the more or less obvious necessity and sufficiency of the R-node constraints.

Lemma 29. *Let μ^\oplus be an R-node of G^\oplus and let X be the set of position variables determined by μ^\oplus together with the PR-ordering variable $\text{ord}(\mu^\oplus)$. A variable assignment α of X can be realized by an embedding of G^\oplus if and only if α satisfies the R-node constraints.*

5.4.2 Relative Positions Determined by P-Nodes

Let μ^\oplus be a P-node of G^\oplus with poles u and v and let $\varepsilon_1, \dots, \varepsilon_k$ be the virtual edges of $\text{skel}(\mu^\oplus)$. Let C be a cycle in G , that induces a cycle κ in $\text{skel}(\mu^\oplus)$. Without loss of generality, κ consists of the virtual edges ε_1 and ε_2 . Let H be a common connected component whose relative position with respect to C is determined by μ^\oplus .

As for R-nodes, we first consider the case that a relative position $\text{pos}_C(H)$ is determined by a cutvertex position variable $\text{pos}_C(v)$. In this case we simply set $\text{pos}_C(H) = \text{pos}_C(v)$. In the following we assume that $\text{pos}_C(H)$ is determined by the P-node μ^\oplus but not by a cutvertex position variable. All cutvertex position variables $\text{pos}_C(v)$ are handled analogously.

As $\text{pos}_C(H)$ is determined by μ^\oplus , the common connected component H is contained in a virtual edge $\varepsilon \in \{\varepsilon_3, \dots, \varepsilon_k\}$ not belonging to κ . Note that embedding ε to the right or to the left of κ determines not only the relative position of H but the position of every connected component that is contained in the expansion graph of ε . We make sure that these relative positions fit to each other by introducing a new variable $\text{pos}_\kappa(\varepsilon)$ with the interpretation that $\text{pos}_\kappa(\varepsilon) = 0$ if and only if ε is embedded to the right of κ . Clearly, $\text{pos}_\kappa(\varepsilon) = \text{pos}_C(H)$ is a necessary condition for every connected component H contained in the expansion graph of ε .

If there is another common cycle C' inducing the same cycle κ in $\text{skel}(\mu)$, we use the same variable $\text{pos}_\kappa(\varepsilon)$ to determine on which side of κ the virtual edge ε is embedded. If C' induces the same cycle oriented in the opposite direction, we use the negation of $\text{pos}_\kappa(\varepsilon)$ instead.

The above constraints are not sufficient for two reasons. First, changing the embedding of $\text{skel}(\mu^\oplus)$ may change edge orderings in the common graph. In this case, there are PR-ordering variables partially determining the embedding of $\text{skel}(\mu^\oplus)$ and we have to make sure that their values and the values of the position variables determined by μ^\oplus fit to each other. Second, if different common cycles induce different cycles in $\text{skel}(\mu^\oplus)$, then not every combination of relative positions with respect to these cycles can actually be achieved by embedding the skeleton.

Connection to Ordering Variables. As mentioned before, we have to add additional constraints to ensure consistency between the position variables and the ordering variables. Let κ be the cycle induced by the cycle C in $\text{skel}(\mu^\oplus)$ and assume without loss of generality that κ consists of the virtual edges ε_1 and ε_2 . Embedding another virtual edge ε to the right or to the left of κ changes the edge ordering at the poles u or v if and only if the expansion graph of ε includes common edges incident to u or v , respectively. In this case, we have a PR-ordering variable determining this embedding choice and we can make sure that the edge ordering and the relative positions fit to each other using an equation or an inequality.

To make this more precise, we define the following five types of virtual edges. Each virtual edge $\varepsilon \in \{\varepsilon_1, \dots, \varepsilon_k\}$ is of exactly one of the following five types.

Type 1. $\text{exp}(\varepsilon)$ includes a common path from u to v .

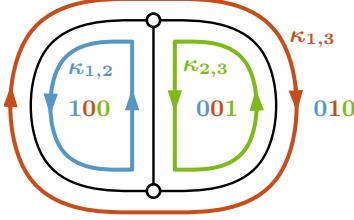


Figure 19. Three virtual edges of Type 1 in a P-node with the three different virtual cycles $\kappa_{1,2}$, $\kappa_{1,3}$, and $\kappa_{2,3}$. For each face, the variable assignment corresponding to this face is given.

Type 2. $\exp(\varepsilon)$ has common edges incident to u and to v but is not of Type 1.

Type 3. $\exp(\varepsilon)$ has a common edge incident to u but none incident to v .

Type 4. $\exp(\varepsilon)$ has a common edge incident to u but none incident to v .

Type 5. $\exp(\varepsilon)$ has no common edges incident to u or to v .

As the cycle κ consists of ε_1 and ε_2 , they must both be of Type 1. Choosing the relative position of ε with respect to κ affects the edge ordering at u or at v if and only if ε is not of Type 5. If ε is of Type 1 or Type 2, then there is a PR-ordering variable $\text{ord}(\mu^\circledcirc)$ determining the order of the edges $\varepsilon_1, \varepsilon_2, \varepsilon$. If the ordering corresponding to $\text{ord}(\mu^\circledcirc) = 0$ has ε to the right of κ , we set $\text{ord}(\mu^\circledcirc) = \text{pos}_\kappa(\varepsilon)$. Otherwise, we set $\text{ord}(\mu^\circledcirc) \neq \text{pos}_\kappa(\varepsilon)$.

If ε is of Type 3, the edge ordering of $\varepsilon_1, \varepsilon_2, \varepsilon$ is determined by the PR-ordering variable $\text{ord}(\mu_s^\circledcirc)$. As before, we get either the equation $\text{ord}(\mu_s^\circledcirc) = \text{pos}_\kappa(\varepsilon)$ or the inequality $\text{ord}(\mu_s^\circledcirc) \neq \text{pos}_\kappa(\varepsilon)$. The case that ε is of Type 4 is analogous, except that we have the PR-ordering variable $\text{ord}(\mu_t^\circledcirc)$ instead of $\text{ord}(\mu_s^\circledcirc)$.

Multiple Cycles. If there are common cycles inducing different cycles in the P-node μ^\circledcirc , then at least three virtual edges must be of Type 1 (i.e., their expansion graph includes a common path between the poles u and v). As we assume that μ^\circledcirc has common P-node degree 3, three virtual edges are of Type 1 and all remaining virtual edges are of Type 5. Let ε_1 , ε_2 , and ε_3 be the virtual edges of Type 1 and let ε be another virtual edge of $\text{skel}(\mu^\circledcirc)$. Denote the cycle consisting of ε_i and ε_j by $\kappa_{i,j}$ ($i, j \in \{1, 2, 3\}$ and $i < j$). To simplify the notation, we use $\text{pos}_{i,j}(\varepsilon)$ as short form for the relative position $\text{pos}_{\kappa_{i,j}}(\varepsilon)$.

Let $\varepsilon \in \{\varepsilon_4, \dots, \varepsilon_k\}$ be another virtual edge of $\text{skel}(\mu^\circledcirc)$. Then we are interested in the three position variables $\text{pos}_{1,2}(\varepsilon)$, $\text{pos}_{1,3}(\varepsilon)$, and $\text{pos}_{2,3}(\varepsilon)$, which are not independent from each other. Moreover, which combinations of relative positions can actually be realized depends on the ordering of ε_1 , ε_2 , and ε_3 . This ordering is determined by the PR-ordering variable $\text{ord}(\mu^\circledcirc)$. In the remainder of this section, we first figure out which combinations of values for $\text{ord}(\mu^\circledcirc)$, $\text{pos}_{1,2}(\varepsilon)$, $\text{pos}_{1,3}(\varepsilon)$, and $\text{pos}_{2,3}(\varepsilon)$ are actually possible and then show that restricting the variables to these combinations is equivalent to a set of equations and inequalities.

When fixing the order variable $\text{ord}(\mu^\circledcirc)$ (without loss of generality to 0), we get the situation shown in Figure 19. From the set of eight combinations for the three position variables, only the three combinations 100, 010 and 001 are possible. When changing the order of the edges (setting $\text{ord}(\mu^\circledcirc)$ to 1), every bit is reversed. Thus, for the tuple $(\text{ord}(\mu^\circledcirc), \text{pos}_{1,2}(\varepsilon), \text{pos}_{1,3}(\varepsilon), \text{pos}_{2,3}(\varepsilon))$ we get the possibilities 0100, 0010, 0001 and their complements 1011, 1101, 1110. A restriction equivalent to this is called *P-node 4-constraint* (where equivalent means that an arbitrary subset of variables may be negated).

We note that, in the short version of this paper [5], we missed the fact that the combinations 1000 and 0111 are not possible. This lead to the wrong assumption that a combination is feasible if and only if there is an odd number of 1s, which can be expressed as a linear equation over \mathbb{F}_2 . As a matter of fact, the P-node 4-constraint allows six different combinations, which is not a power of two and can thus not be the solution space of a system of linear equations over \mathbb{F}_2 . Thus, a P-node 4-constraint is in particular not equivalent to a set of equations or inequalities. We resolve this issue with the following lemma in conjunction with the new results from Section 3.2.

Lemma 30. *If two variables of a P-node 4-constraint are known to be equal or unequal, the P-node 4-constraint is equivalent to a set of equations and inequalities.*

Proof. We basically have the three possibilities 0100, 0010, and 0001 and their complements for the variables $abcd$. No pair of variables is equal in all three possibilities and no pair of variables is unequal in all three possibilities. Thus requiring equality or inequality for one of the pairs eliminates exactly one or two of these three possibilities. If exactly one possibility and its complement remains, this is obviously equivalent to a set of equations and inequalities.

If 0100 and 0010 (and their complements) remain, this is equivalent to $a = d$ and $b \neq c$. If 0100 and 0001 remain, this is equivalent to $a = c$ and $b \neq d$. Finally, if 0010 and 0001 remain, this is equivalent to $a = b$ and $c \neq d$. \square

In the following, we show that for every P-node 4-constraint, we always find an equation or inequality between a pair of variables, turning all P-node 4-constraints into a set of equations and inequalities.

Consider the union graph G^\cup . If the poles $\{u, v\}$ are a separating pair in G^\cup , then each split component is the union of the expansion graphs of several virtual edges of $\text{skel}(\mu^\circledcirc)$. As the expansion graphs of ε_1 , ε_2 , and ε_3 have common uv -paths, we can assume that they are not separated (Theorem 3). Moreover, having common P-node degree 3 implies that none of the other expansion graphs has a common edge incident to u or to v (they are all of Type 5). Thus, again by Theorem 3, we can assume that $\{u, v\}$ is not a separating pair in G^\cup .

It follows that there must be a path π in G^\cup that connects a vertex of $\exp(\varepsilon)$ with a vertex of (without loss of generality) $\exp(\varepsilon_1)$ that does not pass through u or v or vertices of the expansion graphs of ε_2 and ε_3 . It follows that the relative position of ε with respect to $\kappa_{2,3}$ must be the same as the relative position of any internal vertex of $\exp(\varepsilon_1)$ with respect to $\kappa_{2,3}$. As this relative position is determined by the order of ε_1 , ε_2 , and ε_3 , we obtain either $\text{ord}(\mu^\circledcirc) = \text{pos}_{2,3}(\varepsilon)$ or $\text{ord}(\mu^\circledcirc) \neq \text{pos}_{2,3}(\varepsilon)$.

Sufficiency. We call the constraints defined in this section the *P-node constraints*. The following lemma follows directly from the previous considerations.

Lemma 31. *Let μ^\circledast be a P-node of G^\circledast and let X be the set of position variables determined by μ^\circledast together with the PR-ordering variables of μ^\circledast . A variable assignment α of X can be realized by an embedding of G^\circledast if and only if α satisfies the P-node constraints.*

5.4.3 Relative Positions Determined by Common Cutvertices

Recall that the relative position $\text{pos}_C(H)$ is determined by a common cutvertex v of G^\circledast if C contains v and H lies in a split component S^\circledast (with respect to v) different from the split component containing C such that S^\circledast has a common edge incident to v .

First note that the whole split component S^\circledast has to be embedded on one side of C . Thus, for every common connected component in S^\circledast , we would get the same set of constraints. To reduce the number of constraints, we introduce the variable $\text{pos}_C(S^\circledast)$ representing the decision of embedding S^\circledast either to the right or to the left of C . Clearly, $\text{pos}_C(H) = \text{pos}_C(S^\circledast)$ for every common connected component H in S^\circledast is a necessary condition.

Note that this condition is very similar to the first type of constraints we required for P-nodes (connected components in the expansion graph of the same virtual edge have the same relative positions). As for the P-nodes, we have to address two potential issues. First, embedding the split component S^\circledast to one side or another of C changes the edge ordering around the cutvertex v . Second, if there are multiple cycles through v , then the relative positions of S^\circledast with respect to all these cycles must be consistent.

Connection to Ordering Variables. Let B^\circledast be the block of S^\circledast containing v and let e_1 and e_2 be the two edges of C incident to v . Moreover, let e be a common edge of B^\circledast incident to v . Recall that the cyclic order of e_1 , e_2 , and e is described by the cutvertex-ordering variable $\text{ord}(e_1, e_2, B^\circledast)$.

Assume without loss of generality that e_1 is oriented towards v and e_2 is oriented away from v (in the cycle C). Then the (clockwise) cyclic order e_1, e, e_2 forces the block B^\circledast , and thus the whole split component S^\circledast , to lie left of C . The opposite cyclic order forces S^\circledast to the right of C . Thus, depending on the orientation of C , we either get $\text{ord}(e_1, e_2, B^\circledast) = \text{pos}_C(S^\circledast)$ or $\text{ord}(e_1, e_2, B^\circledast) \neq \text{pos}_C(S^\circledast)$ as necessary conditions.

Multiple Cycles. Assume multiple common cycles C_1, \dots, C_k contain the cutvertex v and assume that these cycles are already embedded. We have to make sure that every assignment to the variables $\text{pos}_{C_i}(S^\circledast)$ for $i = 1, \dots, k$ actually corresponds to a face of $C_1 \cup \dots \cup C_k$ that is incident to v .

We cannot directly express this requirement as a set of equations and inequalities on the position variables. However, if we assume that a given variable assignment for the cutvertex-ordering variables of v can be realized by an embedding of G^\circledast (which is ensured by the constraints from Section 5.2), then the above constraints establishing the connection between the cutvertex-ordering variables and the position variables make sure that the corresponding values for the position variables are also realized.

Sufficiency. We call the constraints from this section the *common cutvertex constraints*. Let α be a variable assignment for the cutvertex-ordering variables of a cutvertex v . We say that α is *order realizable*, if G^\oplus admits an embedding realizing α . We obtain the following lemma.

Lemma 32. *Let X be the set of position variables that are determined by the common cutvertex v in G^\oplus and let Y be the cutvertex-ordering variables of v . A variable assignment α of $X \cup Y$ can be realized by an embedding of G^\oplus if and only if α satisfies the common cutvertex constraints and $\alpha|_Y$ is order realizable.*

5.4.4 Relative Positions Determined by Exclusive Cutvertices

As in the previous section, let S^\oplus be the split component with respect to the cutvertex v that contains the connected component H . As before, every common connected component of S^\oplus has to be embedded on the same side of C . However, in this case, we need slightly stronger constraints.

Let H_v be the connected component of the common graph that includes the cutvertex v . Let further $B_1^\cup, \dots, B_k^\cup$ be the union bridges of H_v (note that this is the first time, where the second graph G^\otimes comes into play). As the union bridge B_i^\cup (for $i = 1, \dots, k$) has to be completely embedded into a single face of H_v , every common connected component in B_i^\cup lies on the same side of C . As before for the split components, we represent the decision of putting B_i^\cup to the right or to the left of C using the variable $\text{pos}_C(B_i^\cup)$. Then the constraint $\text{pos}_C(B_i^\cup) = \text{pos}_C(H)$ for every common connected component H in B_i^\cup is clearly necessary. Note that the resulting constraints are strictly stronger than setting $\text{pos}_C(S^\oplus) = \text{pos}_C(H)$ for every common connected component H in S^\oplus , as S^\oplus is contained in a single union bridge.

Recall that (in contrast to the previous section) S^\oplus does not contain a common edge incident to v . It follows that the decision of putting H to the right or to the left of C in an embedding of G^\oplus has no influence on the edge ordering at v . Thus, there is no need for further constraints to ensure consistency between edge orderings and relative positions. Moreover, we will see that there is no need for additional constraints to make sure that the relative positions actually describe a face (in case v is contained in multiple cycles).

Sufficiency. We call the constraints defined in this section the *exclusive cutvertex constraints*. Assume that H_v is already embedded. Then we can choose to embed S^\oplus into an arbitrary face of H_v incident to v , which determines the relative positions of the components in S^\oplus with respect to cycles through v without affecting any other embedding choice. It only remains to make sure that the relative positions of H with respect to all cycles through v actually describe a face incident to v . Unfortunately, the exclusive cutvertex constraints do not guarantee this property and we are not able to give additional constraints enforcing it.

However, we can prove the following lemma by exploiting the fact that $\text{pos}_C(H)$ is in G^\otimes determined by an R-node, by a P-node, or by a common cutvertex. The R-node, P-node, and common cutvertex constraints of G^\otimes then help to prove the existence of the desired face.

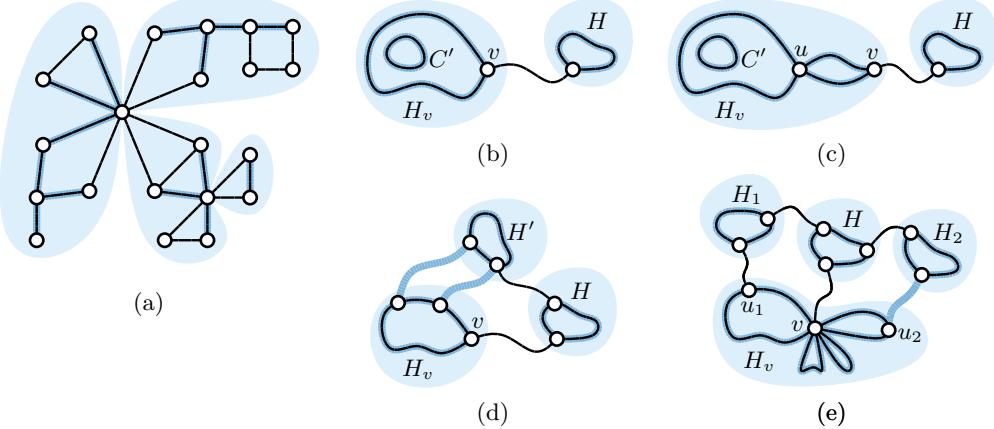


Figure 20. (a) The extended blocks of the graph G^\oplus . (b) The cutvertex v and the cycle C' are contained in the same block of H_v . (c) v and C' are in different blocks. (d) The exclusive bridge containing H' has two attachments in H_v . Note that H' and H belong to the same union bridge. (e) The union bridge containing H (and H_1 and H_2) has attachments in different blocks of H_v , i.e., it is not block-local.

We call the union of all R-node, all P-node, all common cutvertex, and all exclusive cutvertex constraints the *position constraints*.

Lemma 33. *Let G^\oplus and G^\otimes have common P-node degree 3 and simultaneous cutvertices of common degree at most 3. Let α be a variable assignment for the ordering and position variables satisfying the ordering and position constraints with respect to G^\oplus and G^\otimes . Then G^\oplus admits an embedding that realizes α .*

Proof. Let the blocks of G^\oplus be partitioned into a maximum number of partitions such that two blocks that each have a common edge incident to a cutvertex belong to the same partition; see Figure 20a. Let $B_1^\oplus, \dots, B_k^\oplus$ be the blocks of one such partition and let $P^\oplus = B_1^\oplus \cup \dots \cup B_k^\oplus$. Note that P^\oplus is a maximal subgraph of G^\oplus such that every split component with respect to a cutvertex v includes a common edge incident to v . We call P^\oplus an *extended block* of G^\oplus . By Lemmas 27, 29, 31, and 32, P^\oplus (and every other extended block) admits an embedding that realizes α .

Analogously, we define the extended blocks of G^\otimes and choose an embedding for each extended block in G^\otimes that realizes α .

To get an embedding of G^\oplus realizing α , it remains to combine the embeddings of the extended blocks at the cutvertices separating them. Let P_v^\oplus be the extended block of G^\oplus that includes a common edge incident to the cutvertex v and let P^\oplus be another extended block containing v . Let H_v be the connected component of the common graph containing v . Note that H_v is completely contained in P_v^\oplus and thus its embedding is already fixed. Note further that P^\oplus is part of a split component with respect to v and thus part of a single union bridge of H_v . Thus, the exclusive cutvertex constraints make sure that the relative positions with respect to common cycles through v are the same for

all common connected components in P^\oplus . Hence, the embeddings of P_v^\oplus and P^\oplus can be combined such that the resulting embedding of $P_v^\oplus \cup P^\oplus$ realizes α if and only if the relative positions of one common connected component H in the union bridge containing P^\oplus with respect to the common cycles through v describe a face of H_v that is incident to v . The latter follows immediately from the following two claims.

Claim 1. If the relative positions of H with respect to common cycles through v describe a face of H_v , then this face is incident to v .

Claim 2. The relative positions of H with respect to all common cycles through v describe a face of H_v .

We start with the proof of Claim 1. Let f be a face of H_v such that the relative positions of H with respect to cycles through v (as given by α) are the same as the relative positions of f with respect to these cycles. Then f is incident to the cutvertex v for the following reason. If f is not incident to v , then there exists a common cycle C' (not containing v) in the connected component H_v of v that separates v from a connected component H . By Lemma 23 we can assume that C' is part of the extended cycle basis.

There are two possibilities. If v belongs in G^\oplus to the same block as C' (see Figure 20b), the relative position $\text{pos}_{C'}(H)$ is determined by the relative position of v with respect to C' , as v separates H from C' in G^\oplus and C' does not contain v . Thus, we have the equation $\text{pos}_{C'}(v) = \text{pos}_{C'}(H)$ in this case. Otherwise, G^\oplus has another cutvertex u separating v and H from C' ; see Figure 20c. Then v and H are in the same split component with respect to this cutvertex. In this case we also have the requirement that v and H are on the same side of C' . Hence, the cycle C' cannot separate v from H , which proves Claim 1.

It remains to prove Claim 2. Let B^\cup be the union bridge of H_v that contains H . Recall that the exclusive cutvertex constraints require $\text{pos}_C(H') = \text{pos}_C(B^\cup)$ for every cycle C of H_v and every common connected component H' of B^\cup . Thus, showing that the relative positions of H' describe a face of H_v for one common connected component of B^\cup shows this fact for all common connected components of B^\cup (and thus in particular for H).

By Theorem 6, we can assume one of the following is true. The common connected component H_v is a cycle; the union bridge B^\cup is not block-local; or B^\cup is not exclusive one-attached.

If H_v is a cycle, then there is only a single cycle through v and both sides of this cycle form a face of H_v . Thus, there is nothing to show in this case.

Assume the union bridge B^\cup is not exclusive one-attached. Then there exists without loss of generality a ②-bridge $B^②$ that belongs to the union bridge B^\cup such that $B^②$ has two attachments in H_v ; see Figure 20d (the case of a ①-bridge $B^①$ belonging to B^\cup and having two attachments in H_v is analogous). Let further H' be a common connected component contained in $B^②$. Then H' belongs in $G^②$ to a block that contains at least one block of H_v . Thus, the extended block containing H' completely contains H_v . It follows that the relative positions of H' with respect to cycles in H_v describe a face of H_v .

Finally, assume B^\cup is not block-local. Then there are two ①-bridges B_1^\oplus and B_2^\oplus (with $i \in \{1, 2\}$) belonging to the union bridge B^\cup with attachments in different blocks of H_v . If one of these bridges has an attachment vertex in a block of H_v not containing the

cutvertex v , then the relative positions of this bridge with respect to any common cycle containing v is determined by an R-node, by a P-node, or by a common cutvertex. Thus, the relative positions correspond to a face of H_v in this case. It remains to consider the case that the attachment vertices of B_1^\oplus and B_2^\oplus belong to blocks of H_v incident to v ; see Figure 20e.

Let S_1, \dots, S_k be the split components of the common component H_v with respect to the cutvertex v . Assume without loss of generality that B_1^\oplus and B_2^\oplus have their attachment vertices u_1 and u_2 in S_1 and S_2 , respectively. Let H_1 and H_2 be common connected components in B_1^\oplus and B_2^\oplus , respectively. The relative position of H_1 with respect to a cycle through v that is not contained in S_1 is determined (in G^\oplus) by the common cutvertex v . Thus, the relative positions of H_1 with respect to $S_2 \cup \dots \cup S_k$ describe a face of $S_2 \cup \dots \cup S_k$. Moreover, this face contains the whole split component S_1 . Thus, if the relative positions of H_1 with respect to cycles in S_1 describe a face of S_1 , then the relative positions with respect to cycles in $H_v = S_1 \cup \dots \cup S_k$ describe a face of H_v . Clearly, this is true as H_1 and H_2 have the same relative positions (they are in the same union bridge B^\cup) and the relative positions of H_2 with respect to cycles in S_1 describe a face of S_1 (one can use a symmetric argument to the one above). This concludes the proof. \square

5.4.5 Computing the Constraints

Recall from Section 5.2 (Lemma 27) that we have potentially $O(n^3)$ cutvertex-ordering variables. Moreover, there are $O(n^2)$ cycles in the extended cycle basis \mathcal{C} and thus $O(n^3)$ component position variables. Thus, our aim is to compute the position constraints described in the previous sections in $O(n^3)$ time.

Let $C \in \mathcal{C}$ be a cycle. For the relative positions with respect to C that are determined by R-nodes or P-nodes, we need to know for every R- and P-node μ of G^\oplus and of G^\otimes , whether it induces a cycle κ in $\text{skel}(\mu)$. If so, we also need to know the cycle κ . This can clearly be done in linear time for each cycle, yielding a total running time of $O(n^3)$ (note that techniques from [7] can be used to compute this information for multiple cycles in linear time).

Similarly, in $O(n^2)$ time, we can compute for every virtual edge ε in $\text{skel}(\mu)$, which common connected components are contained in the expansion graph of ε . Assume μ is an R-node and let X be the set of ordering variables determined by μ . With the above information, one can easily compute the R-node constraints for μ in $O(|X| + |\text{skel}(\mu)|)$ time. As each relative position is determined by at most one R-node, the sets X are disjoint for different R-nodes. Thus, we get a total running time of $O(n^3)$ for computing the R-node constraints.

Computing the P-node constraints of a P-node μ can be done analogously (yielding $O(n^3)$ running time in total), except for the case where we have to handle a P-node 4-constraint. Recall that we get P-node 4-constraints if three virtual edges $\varepsilon_1, \varepsilon_2$, and ε_3 of $\text{skel}(\mu)$ include common paths between the poles. For every other virtual edge ε , we then get a P-node 4-constraint, which makes $O(|\text{skel}(\mu)|)$ P-node 4-constraints for μ . For the P-node 4-constraint corresponding to the virtual edge ε , we have to check

whether the union graph G^{\cup} has a path π from $\exp(\varepsilon)$ to $\exp(\varepsilon_i)$ (for $i \in \{1, 2, 3\}$) that is disjoint from the expansion graphs $\exp(\varepsilon_j)$ for $j \in \{1, 2, 3\}$ with $i \neq j$. This can clearly be done in $O(n)$ time for each edge ε of $\text{skel}(\mu)$. It follows that we can compute the P-node constraints in $O(n^3)$ time.

For a cutvertex v of G^{\oplus} , consider the relative positions X determined by the common cutvertex v . For every split component S^{\oplus} and every common connected component H in S^{\oplus} , we have the constraint $\text{pos}_C(H) = \text{pos}_C(S^{\oplus})$. These constraints can be easily computed in $O(n + |X|)$ time. As the sets X are disjoint for every cutvertex, this yields a total running time of $O(n^3)$. Moreover, we have to compute constraints of the type $\text{ord}(e_1, e_2, B^{\oplus}) = \text{pos}_C(S^{\oplus})$ connecting the relative positions to the cutvertex ordering variables. Clearly, for each variables $\text{pos}_C(S^{\oplus})$ this constraint can be added in constant time, which yields a running time in $O(|X|)$. Hence, the common cutvertex constraints can be computed in $O(n^3)$ time.

Finally, consider the relative positions X determined by the exclusive cutvertex v and let H_v be the common connected component containing v . For every union bridge B^{\cup} , every common connected component H in B^{\cup} , and every common cycle through v , we have to add the constraint $\text{pos}_C(B^{\cup}) = \text{pos}_C(H)$. We can first (in $O(n)$ time) partition the common connected components according to their union bridges. Then, adding these constraints for one cycle C can be done in $O(|X_C|)$ time, where X_C is the set of relative positions with respect to C in X . Thus, we get the exclusive cutvertex constraints for v in $O(|X|)$ time, which yields a total running time of $O(n^3)$.

Lemma 34. *The position constraints can be computed in $O(n^3)$ time.*

5.5 Putting Things Together

Assume $(G^{\oplus}, G^{\otimes})$ is a SEFE instances such that G^{\oplus} and G^{\otimes} are connected graphs, G^{\oplus} and G^{\otimes} have common P-node degree 3, and every simultaneous cutvertex has common degree 3.

We first used Theorem 1 to get rid of all union cutvertices. This helped to ensure consistent edge orderings in Section 5.2. Actually, without union cutvertices, we know for each vertex v that it is either not a cutvertex in one of the graphs G^{\oplus} or G^{\otimes} , which makes representing the possible edge orderings much simpler, or that it has common degree 3, which also makes the ordering simple.

To ensure consistent relative positions of two common connected components with respect to each other, we first showed in Section 5.1 that it suffices to ensure consistent relative positions of each common connected component with respect to the cycles of a cycle basis in the other component. Unfortunately, setting relative positions with respect to cycles does not necessarily lead to an embedding (e.g., if a cycle C_1 lies “inside” C_2 , and C_2 lies “inside” C_3 , then C_3 cannot lie “inside” C_1).

This leads to difficulties, when one component H_1 can be potentially embedded into several faces of another component H_2 , which is the case when H_1 is attached to H_2 via only two vertices that are a separating pair of H_2 , or when H_1 and H_2 are separated by a cutvertex. For the former case, it helped to assume that split components of union

separating pairs have a very special structure (Theorem 3). For the latter case, it helped to assume that there are no union bridges that are block-local and exclusively one-attached (Theorem 6).

Using Theorem 6 comes at the cost that we have to satisfy some common-face constraints. However, in Section 5.3 we showed that this can be done easily (Lemma 28).

The set of equations and inequalities we obtain has size $O(n^3)$, can be computed in $O(n^3)$ time, and can be solved in linear time in its size [2]. This lets us conclude with the following theorem.

Theorem 9. *SEFE can be solved in $O(n^3)$ time for two connected graphs with common P-node degree 3 and simultaneous cutvertices of common degree at most 3.*

6 Conclusion

In this paper, we have shown how to combine techniques for ensuring consistent relative positions [7] with known [1, 4] and newly developed tools for ensuring consistent edge orderings. This lead to an efficient algorithm solving SEFE for two connected graphs with common P-node degree 3 and simultaneous cutvertices of common degree at most 3. Together with the linear-time algorithm for decomposing a given instance into equivalent instances in which each 2-component is a cycle, this gives an efficient algorithm if each connected component of the common graph is biconnected, has maximum degree 3, or is outerplanar with maximum degree 3 cutvertices.

We note that all techniques developed in Section 5 extend to the sunflower case, where we have multiple graphs pairwise intersecting in the same common graph. Actually, the two graphs G^\oplus and G^\otimes are only considered together if G^\otimes restricts the embedding choices of the common graph in G^\oplus in a way that makes it possible to formulate certain constraints. Thus, more graphs intersecting in the same common graph can only help. Moreover, the preprocessing algorithms from Section 3 also directly extend to the sunflower case when adapting the definition of impossible P-nodes (Lemma 5) in a straightforward manner.

Besides solving this fairly general set of SEFE instances, our results and in particular the preprocessing algorithms give some new structural insights that may help in further research. For example, Theorem 2, stating that one can assume all union-link graphs to be connected, not only helps in later chapters but also shows that the decision of ordering virtual edges in P-nodes of the common graph is fairly easy.

What remains poorly understood are the edge orderings at cutvertices. We were basically able to handle cutvertices if the choices boil down to binary decision. This is for example the case if the cutvertex has only common degree 3. Although less obvious, this is also the case if the instance has common P-node degree 3. For a cutvertex in G^\oplus , this basically means that the other graph G^\otimes hierarchically groups the common edges incident to the cutvertex such that there are at most three groups on each level, yielding binary decisions.

We feel that the understanding of edge orderings at cutvertices is a major bottleneck for solving more general instances of SEFE. To get a better understanding, we believe

that it would be helpful to first study cutvertices in the context of constrained planarity problems, which are often a bit more simple. A good starting point could be the constrained planarity problems planarity with partitioned PQ-constraints or variants like partitioned full-constraints [6].

References

- [1] Patrizio Angelini, Giuseppe Di Battista, Fabrizio Frati, Maurizio Patrignani, and Ignaz Rutter. Testing the simultaneous embeddability of two graphs whose intersection is a biconnected or a connected graph. *Journal of Discrete Algorithms*, 14:150–172, 2012.
- [2] Bengt Aspvall, Michael F. Plass, and Robert Endre Tarjan. A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Information Processing Letters*, 8(3):121–123, 1979.
- [3] Thomas Bläsius, Stephen G. Kobourov, and Ignaz Rutter. *Handbook of Graph Drawing and Visualization*, chapter Simultaneous Embedding of Planar Graphs, pages 349–381. Chapman and Hall/CRC, 2013.
- [4] Thomas Bläsius and Ignaz Rutter. Simultaneous PQ-ordering with applications to constrained embedding problems. In *Proc. 24th ACM-SIAM Sympos. Discrete Algorithm*, SODA’13, pages 1030–1043. ACM, 2013.
- [5] Thomas Bläsius, Annette Karrer, and Ignaz Rutter. Simultaneous embedding: Edge orderings, relative positions, cutvertices. In Stephen Wismath and Alexander Wolff, editors, *Graph Drawing (GD’13)*, volume 8242 of *LNCS*, pages 220–231. Springer, 2013.
- [6] Thomas Bläsius and Ignaz Rutter. A new perspective on clustered planarity as a combinatorial embedding problem. In Christian Duncan and Antonios Symvonis, editors, *Graph Drawing (GD’14)*, volume 8871 of *LNCS*, pages 440–451. Springer, 2014.
- [7] Thomas Bläsius and Ignaz Rutter. Disconnectivity and relative positions in simultaneous embeddings. *Computational Geometry: Theory and Applications*, 48(6):459–478, 2015.
- [8] Kellogg S. Booth and George S. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. *Journal of Computer and System Sciences*, 13:335–379, 1976.
- [9] G. Di Battista and R. Tamassia. On-line maintenance of triconnected components with SPQR-trees. *Algorithmica*, 15(4):302–318, 1996.

- [10] Harold N. Gabow and Robert Endre Tarjan. A linear-time algorithm for a special case of disjoint set union. *Journal of Computer and System Sciences*, 30(2):209–221, 1985.
- [11] Elisabeth Gassner, Michael Jünger, Merijam Percan, Marcus Schaefer, and Michael Schulz. Simultaneous graph embeddings with fixed edges. In FedorV. Fomin, editor, *Graph-Theoretic Concepts in Computer Science (WG'06)*, volume 4271 of *LNCS*, pages 325–335. Springer, 2006.
- [12] Carsten Gutwenger and Petra Mutzel. A linear time implementation of SPQR-trees. In Joe Marks, editor, *Graph Drawing (GD'00)*, volume 1984 of *LNCS*, pages 77–90. Springer, 2001.
- [13] Bernhard Haeupler, Krishnam Jampani, and Anna Lubiw. Testing simultaneous planarity when the common graph is 2-connected. *J. Graph Algorithms Appl.*, 17(3):147–171, 2013.
- [14] Dov Harel and Robert Endre Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM Journal on Computing*, 13(2):338–355, 1984.
- [15] Michael Jünger and Michael Schulz. Intersection graphs in simultaneous embedding with fixed edges. *Journal of Graph Algorithms and Applications*, 13(2):205–218, 2009.
- [16] Łukasz Kowalik and Maciej Kurowski. Short path queries in planar graphs in constant time. In *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing STOC '03*, pages 143–148. ACM, 2003.
- [17] Marcus Schaefer. Toward a theory of planarity: Hanani-tutte and planarity variants. In Walter Didimo and Maurizio Patrignani, editors, *Graph Drawing (GD'12)*, volume 7704 of *LNCS*, pages 162–173. Springer, 2013.